PERFORMATIVE MEAN FIELD GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper investigates performative prediction in large-population games where the number of players is unknown and effectively infinite. In many real-world settings, such as traders adjusting investment strategies or merchants setting prices, any single player's decision has negligible influence on the overall environment, yet the collective decision distribution of all players significantly shapes the data distributions that each player observes and makes decisions on. Existing performative prediction frameworks, which rely on discrete, finite-player models, fail to capture these dynamics, since a change in any single player's discrete decision does not meaningfully alter the population-level decision distribution and thus induces no perceptible shift in the induced data distributions. We address this challenge by introducing the *performative mean field game* framework. In this framework, a continuum of players interacts through an unknown mapping from the population decision distribution to the resulting data distributions. We define the *performatively stable mean-field equilibrium* as the fixed point where the decision distribution is exactly the aggregate of each player's optimal response to the data distribution induced by that same decision distribution, eliminating the need for further decision updates. Our main technical contributions are threefold. First, we formalize the above game and equilibrium concept. Second, we introduce a sensitivity measure in Sobolev norms that quantifies the impact of distribution shifts on individual decision-making and use it to establish theoretical convergence conditions for performatively stable mean-field equilibria. Third, building on these theoretical results, we design an adaptive algorithm that estimates sensitivity on the fly and provably converges to a performatively stable mean-field equilibrium without requiring prior knowledge of the underlying decision-to-data mapping. Experiments on quantitative trading, recommendation system, and ride-sharing market validate our theoretical results and demonstrate the effectiveness of the proposed algorithm.

1 INTRODUCTION

This paper studies games with an unknown and effectively infinite number of players, where the collective decision distribution of all players induces shifts in the data distributions under which each player optimizes. For instance, when all merchants in a product category raise prices to maximize their profits, the category's number of clicks declines and the total revenue drops, yet any individual merchant's price change cannot change the number of clicks for the category. In such practical scenarios, the data distributions are shaped by the decision distribution, while any single player's decision has negligible influence, such as financial markets with millions of trading accounts, online marketplaces with thousands of merchants, and ride share market with many consumers.

The phenomenon of decision-dependent distribution shifts has been formalized as *performative prediction* (Perdomo et al., 2020) and widely studied across applications, including reinforcement learning (Pollatos et al., 2025), fairness (Jin et al., 2026), data streams (Gower-Winter et al., 2025), and federated learning (Jin et al., 2024; Qian et al., 2026). Performative prediction was introduced for single-agent settings (Perdomo et al., 2020) and extended to multiplayer games (Narang et al., 2023), where a data distribution map $\mathcal{D}_i(\cdot)$ maps the joint decision $X = (x_1, \ldots, x_n)$ of n players to player i's data distribution $\mathcal{D}_i(X)$. Both the single-agent and multiplayer frameworks focus on settings where individual decisions meaningfully influence the data distributions, which is natural in applications such as pricing decisions of a few large firms or national policy-making. However,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

in large-population settings, only the collective decision distribution meaningfully shapes the data distributions rather than any single individual decision. The discrete models in existing performative prediction frameworks model the data distribution maps as a function of individual decisions, and thus fail to capture the decision-distribution-dependent dynamic in such scenarios. Moreover, the number of players is unknown and effectively infinite in these scenarios, which creates two computational challenges for existing frameworks. First, the data distribution map $\mathcal{D}_i(x_1, \ldots, x_n)$ cannot be formulated without knowing the number of players n. Second, when n is large, solving for an equilibrium becomes computationally intractable due to the high dimensionality of the joint decision space. Consequently, existing performative prediction frameworks are inapplicable in these scenarios, motivating a different modeling framework.

To study performative prediction in scenarios with an unknown and effectively infinite number of players, we model the players as a continuum (Aumann, 1964) and adopt mean field game theory (Lasry & Lions, 2007), where each player determines the optimal decision with respect to the collective decision distribution rather than individual decisions. In our framework, a player characterized by a type $\theta \in \Theta$ optimizes its decision $x \in \mathcal{X}$ based on the type-decision distribution $m \in \mathcal{P}(\Theta \times \mathcal{X})$ and a data distribution $\mathcal{D}_\theta(m)$ induced by the decision distribution m. This formulation is independent of the specific number of players and avoids jointly solving best-response conditions of all players to compute an equilibrium. We term this game a *Performative Mean Field Game*. Here, we illustrate our framework through a performative recommendation system game.

Example. *(Recommendation System) Each online merchant with cost $\theta \in \Theta$ chooses a price adjustment $x \in \mathcal{X}$ to minimize its expected loss: $\min_x \mathbb{E}_{Z\sim\mathcal{D}(m)}\Big[-(1+x)(p-\theta)\cdot Z + (1+x)\left(1+\int_{\Theta\times\mathcal{X}} x' m(d\theta, dx')\right)\Big]$, where p is the price of this product category and $Z \sim \mathcal{D}(m)$ is its number of clicks. The first term is the negative profit from clicks. The second term is the congestion cost, capturing the reduction in sales when the merchant's own price and the average price are high. A reasonable data distribution map is*

$$Z \sim \mathcal{D}(m) = Z_0 \left(1 - \kappa \left(\int_{\Theta\times\mathcal{X}} x' m(d\theta, dx')\right)^3\right) + w, \tag{1}$$

where the cubic form captures the nonlinear shifts in the number of clicks induced by average price adjustments $\int_{\Theta\times\mathcal{X}} x' m(d\theta, dx')$, Z_0 is the baseline number of clicks, κ is the performative effect strength, and w is a noise term. Here, the data distribution map is the same for all types of players, i.e., $\mathcal{D}_\theta(m) = \mathcal{D}(m)$.

Given that the data distribution shifts with the decision distribution, a player's current decision may no longer be optimal under the induced data distribution. In response, players repeatedly compute a mean-field equilibrium under the current data distributions, which induces new data distributions, necessitating further decision updates. We term this procedure the *Repeated Performative Mean Field Game*, where at each step $t+1$, each player optimizes its decision $x^{(t+1)}$ under data distribution $\mathcal{D}_\theta(m^{(t)})$ induced by the previous decision distribution $m^{(t)}$. A desirable outcome is a decision distribution m^{PS} that is composed of the optimal decisions of all players with respect to m^{PS} under the data distributions $\mathcal{D}_\theta(m^{PS}), \forall\theta \in \Theta$ induced by m^{PS}, eliminating the need for further decision updates. In this paper, we term such a fixed point a *Performatively Stable Mean-Field Equilibrium*.

1.1 Contributions

We propose a framework for performative mean field games, making the following contributions:

1. **Formalization of Performative Mean Field Games:** In Section 2, we formalize game-theoretic definitions of the performative mean field game. We then describe the procedure of the repeated performative mean field game, where all players repeatedly optimize decisions under induced data distributions. Finally, we define the performatively stable mean-field equilibrium, which is a decision distribution composed of the optimal decisions of all players under induced data distributions.
2. **Equilibrium Convergence Analysis:** To characterize the performative effect, we define an ε-sensitivity measure for the repeated performative mean field game in Section 3.1, which

quantifies the difference in players' losses under distinct induced distributions. Leveraging this measure, we derive theoretical conditions under which a repeated performative mean field game converges to a performatively stable mean-field equilibrium in Section 3.2. We extend this guarantee to the finite-sample setting in Section 3.3, where players can only access empirical loss functions.

3. **Algorithm Design:** Leveraging our theoretical results, we develop an algorithm in Section 4 for repeated performative mean field games without requiring prior knowledge of the data distribution maps. This algorithm ensures convergence to a performatively stable mean-field equilibrium by enforcing the convergence conditions to be satisfied during a repeated performative mean field game.

We evaluate the theoretical results and the proposed algorithm through games in quantitative trading, recommendation system, and ride share market in Section 5. Experimental results with various distribution maps validate our theoretical results and demonstrate the effectiveness of our algorithm.

1.2 Related Work

Performative Prediction. The phenomenon of decisions inducing shifts in data distributions is formalized as *performative prediction* (Perdomo et al., 2020). Research in single-agent and multi-player settings has focused on two objectives: performative stability and performative optimality. Performative stability requires that a decision (or a joint decision of all players) is optimal (or an equilibrium) under the data distribution(s) it induces, via repeated retraining (Perdomo et al., 2020; Ray et al., 2022; Zhong et al., 2026b;a) or repeated gradient descent (Yan & Cao, 2024; Li et al., 2025; Le Cadre et al., 2025). Performative optimality seeks a decision (or a joint decision) that minimizes the loss across all possible induced distributions by incorporating the learned data distribution map into the optimization (Miller et al., 2021; Narang et al., 2023; Tsoy et al., 2025). Despite their differences, these frameworks focus on settings where individual decisions meaningfully influence the data distributions. A distinct scenario arises in large-population settings, where any single decision's impact is negligible and only the collective decision distribution shapes the data distributions.

Mean Field Games. Since our work leverages the mean field game framework, we provide a brief review of continuum of players modeling and mean field game research. Modeling large populations as a continuum of players (Aumann, 1964) enables modeling interactions where a single decision is negligible. The mean field game (Lasry & Lions, 2007), in which each player makes optimal decisions with respect to the collective decision distribution, avoids jointly solving the best-response conditions of all players. Recent research on learning mean field games (Guo et al., 2019; Xie et al., 2021; Cui et al., 2024; Zhang et al., 2025) studies equilibrium computation under transition functions that depend on the current state and decision distributions, either by learning the transition dynamics or by repeatedly sampling losses, aiming to minimize cumulative loss across time steps. In this paper, we investigate the performative mean field game. Since the data distributions depend only on the deployed m, a performatively stable mean-field equilibrium is a fixed point defined as the mean-field equilibrium on the induced data distributions through the unknown data distribution maps $\mathcal{D}_\theta(m)$. We establish convergence conditions to a performatively stable mean-field equilibrium under a natural dynamic, in which every player selects its best response at every iteration.

2 Problem Statement

2.1 Performative Mean Field Game

In a performative mean field game, players are modeled as a continuum and seek to optimize decisions given the type-decision distribution m under induced data distributions. A player with type θ aims to minimize the loss under data distribution $\mathcal{D}_\theta(m)$ induced by type-decision distribution m:

$$\min_{x\in\mathcal{X}} \mathbb{E}_{Z_\theta\sim\mathcal{D}_\theta(m)}\ell(x, Z_\theta, \theta, m), \tag{2}$$

where $\theta \in \Theta \subset \mathbb{R}^{d_\Theta}$ is the type of a player with Θ denoting the nonempty compact type space with Lipschitz boundary, $x \in \mathcal{X} \subset \mathbb{R}^{d_\mathcal{X}}$ is a decision of a player with $\mathcal{X}$ denoting the nonempty compact decision space with Lipschitz boundary, $m \in \mathcal{P}(\Theta \times \mathcal{X})$ is the type-decision distribution on $\Theta \times \mathcal{X}$ with $\mathcal{P}(\Theta \times \mathcal{X})$ denoting the set of all probability measures, $\mathcal{D}_\theta(m)$ is the data distribution induced

by m of a player with type θ, $Z_\theta \sim \mathcal{D}_\theta(m)$ is an instance drawn from the data distribution $\mathcal{D}_\theta(m)$ induced by m, $\ell(\cdot)$ is the loss function, and $d_{\Theta\times\mathcal{X}} = d_\Theta + d_\mathcal{X}$ is the dimension of the joint space $\Theta \times \mathcal{X}$. In this paper, we consider a fair game where the decision space $\mathcal{X}$ is the same for all players. The type distribution $\mu \in \mathcal{P}(\Theta)$ is given and fixed. Every type-decision distribution m decomposes as $m(d\theta, dx) = \mu(d\theta)m(dx|\theta)$, where $m(dx|\theta)$ is the conditional decision distribution given θ. We denote the subset of $\mathcal{P}(\Theta \times \mathcal{X})$ whose marginal distribution on Θ is μ as $\mathcal{P}_\mu(\Theta \times \mathcal{X})$.

For clarity, for a player with type θ, we denote the loss evaluated on m under the data distribution induced by $\bar{m}$ as:

$$\mathcal{L}_\theta(\bar{m}, m) := \mathbb{E}_{Z_\theta \sim \mathcal{D}_\theta(\bar{m})}\ell(x, Z_\theta, \theta, m), \tag{3}$$

where m is the type-decision distribution evaluated in the loss function, and $\bar{m}$ is the type-decision distribution that induces the data distribution.

Given the data distribution $\mathcal{D}_\theta(\bar{m})$ induced by the type-decision distribution $\bar{m}$, the mean-field equilibrium m^* consists of the decisions that minimize the loss function of each player for all types $\theta \in \Theta$, which can be defined as:

$$m^*\left(\left\{(\theta, x) \in \Theta \times \mathcal{X} : \mathcal{L}_\theta(\bar{m}, m^*) = \min_{x\in\mathcal{X}} \mathbb{E}_{Z_\theta \sim \mathcal{D}_\theta(\bar{m})}\ell(x, Z_\theta, \theta, m^*)\right\}\right) = 1, \tag{4}$$

where the decision x of the player with type θ is optimal under the data distribution $\mathcal{D}_\theta(\bar{m})$ induced by the type-decision distribution $\bar{m}$. The optimal decisions x of all players with respect to m^* constitute the type-decision distribution m^*. We denote such a mean-field equilibrium under the mean field game induced by type-decision distribution $\bar{m}$ as $MFE(\bar{m})$:

$$MFE(\bar{m}) := \left\{m^* : m^* \text{ is a mean-field equilibrium of the game induced by } \bar{m}\right\}. \tag{5}$$

2.2 Repeated Performative Mean Field Game

Deploying players' mean-field equilibrium decisions changes the data distributions under which they optimize, rendering these decisions suboptimal. Without knowledge of the underlying distribution maps, players repeatedly update their decisions to minimize loss over the evolving data distributions. We formalize these dynamics as the repeated performative mean field game:

Definition 1 (Repeated Performative Mean Field Game)**.** *Repeated performative mean field game refers to the procedure where, at step $t+1$, the mean-field equilibrium $m^{(t+1)}$ is determined based on data distributions $\mathcal{D}_\theta(m^{(t)})$ for all $\theta \in \Theta$, induced by the type-decision distribution $m^{(t)}$:*

$$m^{(t+1)} \in MFE(m^{(t)}), \tag{6}$$

where $m^{(0)}$ is an initial type-decision distribution inducing the data distributions. We denote $\mathcal{M}_{RPMFG}$ as the closure of the set of type-decision distributions that arise as mean-field equilibria in a repeated performative mean field game. Consequently, every $m^{(t)}$ generated by a repeated performative mean field game belongs to $\mathcal{M}_{RPMFG}$, i.e., $m^{(0)}, m^{(1)}, \cdots, m^{(t)}, \cdots \in \mathcal{M}_{RPMFG}$.

Note that $\mathcal{M}_{RPMFG}$ is defined without circularity and is determined by the performative mean field game structure, so that the distributional form of the equilibria of the game can be characterized without running the repeated game. By focusing on $\mathcal{M}_{RPMFG}$, we only need to handle the specific distributional forms that arise as mean-field equilibria in a given game, rather than arbitrary probability measures on $\Theta \times \mathcal{X}$. For instance, if the game only admits Dirac mean-field equilibria, $\mathcal{M}_{RPMFG}$ consists solely of Dirac measures.

2.3 Performatively Stable Mean-Field Equilibrium

In the repeated performative mean field game, a desirable outcome is a type-decision distribution that is a mean-field equilibrium under the data distributions induced by this mean-field equilibrium, eliminating the need for further decision updates. We formalize this outcome as the performatively stable mean-field equilibrium:

Definition 2 (Performatively Stable Mean-Field Equilibrium)**.** *A type-decision distribution $m^{PS} \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$ is a performatively stable mean-field equilibrium if it is a mean-field equilibrium on data distributions induced by m^{PS}:*

$$m^{PS} \in MFE(m^{PS}). \tag{7}$$

3 THEORETICAL GUARANTEES FOR PERFORMATIVE MEAN FIELD GAME

3.1 ε-SENSITIVITY

To analyze the dynamics of a repeated performative mean field game, we quantify the impact of data distribution shifts on individual decision-making through the ε-sensitivity.

Definition 3 (ε-Sensitivity)**.** *The data distribution maps $\mathcal{D}_\theta(m)$ are ε-sensitive if, for all type-decision distributions $m, m_1, m_2 \in \mathcal{M}_{RPMFG}$, with the loss function $\ell \in H^s(\Theta \times \mathcal{X})$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, the following inequality holds during a repeated performative mean field game:*

$$\|\mathcal{L}_\theta(m_1, m) - \mathcal{L}_\theta(m_2, m)\|_{H^s(\Theta\times\mathcal{X})} \leq \varepsilon \|m_1 - m_2\|_{H^{-s}(\Theta\times\mathcal{X})}, \tag{8}$$

where $\|\cdot\|_{H^s(\Theta\times\mathcal{X})}$ is the norm in Sobolev space $H^s(\Theta \times \mathcal{X})$, $H^s(\Theta \times \mathcal{X})$ consists of all real-valued square-integrable functions on $\Theta \times \mathcal{X}$ whose weak derivatives up to order s are also square-integrable, $\|\cdot\|_{H^{-s}(\Theta\times\mathcal{X})}$ is the negative-order Sobolev norm in the dual space $H^{-s}(\Theta \times \mathcal{X})$.

This ε-sensitivity measure is motivated by the observation that shifts in data distributions directly lead to changes in the players' losses. Unlike prior work (Perdomo et al., 2020; Narang et al., 2023) that measures the data distribution shift via distributional distances (e.g., Wasserstein-1), our measure captures the impact of distribution shifts on decision-making. On the type-decision distribution side, we use the H^{-s} Sobolev norm (Gomes & Mitake, 2015) to quantify differences between type-decision distributions, which can characterize distances between arbitrary distributions (e.g., Dirac, softmax, or Gaussian densities). By only considering type-decision distributions in $\mathcal{M}_{RPMFG}$, we avoid computing norms across arbitrary probability measures. For example, when the equilibrium is a Dirac measure, the H^{-s} norm is evaluated only among Dirac measures. An approach for estimating ε is provided in Algorithm 1 (Section 4), where the norms $\|\cdot\|_{H^s(\Theta\times\mathcal{X})}$ and $\|\cdot\|_{H^{-s}(\Theta\times\mathcal{X})}$ can be computed via the Fourier transform (Cerdà, 2010).

3.2 THEORETICAL GUARANTEE FOR PERFORMATIVELY STABLE MEAN-FIELD EQUILIBRIUM

We first introduce the α-strong monotonicity condition, which is extended from the Lasry-Lions monotonicity (Lasry & Lions, 2007):

Definition 4 (α-Strong Monotonicity)**.** *A mean field game is α-strongly monotone with $\alpha > 0$ if, for all type-decision distributions $\bar{m}, m_1, m_2 \in \mathcal{M}_{RPMFG}$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, the following holds:*

$$\alpha\|m_1 - m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})} \leq \int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) - \mathcal{L}_\theta(\bar{m}, m_2)\right)(m_1 - m_2)(d\theta, dx). \tag{9}$$

Many classes of mean field games, including linear-quadratic games (Carmona & Delarue, 2018) and congestion models (Gomes & Mitake, 2015), satisfy Lasry-Lions monotonicity, corresponding to $\alpha = 0$ in Definition 4. In Sections 4 and 5, we demonstrate that congestion regularization ensures monotone games satisfy α-strong monotonicity, guaranteeing the applicability of our framework.

Together with the compactness of $\Theta, \mathcal{X}$, α-strong monotonicity guarantees the existence and uniqueness of a mean-field equilibrium in each iteration of a repeated performative mean field game, thus ensuring tractable game dynamics and facilitating analysis of the performative effect. We formalize this property in Lemma 1 (Appendix B), which adapts standard proofs (Lacker, 2018) to our setting.

We can now present the dynamics (Theorem 1) and the convergence conditions (Theorem 2).

Theorem 1. *If the performative mean field game is α-strongly monotone, the data distribution maps are ε-sensitive for all $\theta \in \Theta$, the loss $\ell(\cdot)$ is jointly continuous on the product space $\Theta \times \mathcal{X} \times \mathcal{P}(\Theta \times \mathcal{X})$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, then the repeated performative mean field game in Definition 1 produces a sequence of type-decision distributions $\{m^{(0)}, m^{(1)}, \cdots, m^{(t)}, \cdots\} \in \mathcal{M}_{RPMFG}$ such that:*

$$\left\|m^{(t+1)} - m^{(t)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \frac{\varepsilon}{\alpha}\left\|m^{(t)} - m^{(t-1)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \left(\frac{\varepsilon}{\alpha}\right)^t \left\|m^{(1)} - m^{(0)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{10}$$

Theorem 2. *If the conditions of Theorem 1 hold and $\alpha > \varepsilon$, then the repeated performative mean field game in Definition 1 converges linearly to a performatively stable mean-field equilibrium $m^{PS} \in \mathcal{M}_{RPMFG}$ with rate $\frac{\varepsilon}{\alpha}$.*

The proofs of Theorems 1 and 2 are given in Appendices C and D. Theorem 2 assumes full access to the true data distributions, which may not hold in practice. To enhance its applicability, we extend Theorem 2 to the finite-sample setting in the following section.

3.3 Finite Sample Analysis

In this section, we establish the convergence conditions when players can only collect finite samples from data distributions. We adopt the following standard finite variance condition (Narang et al., 2023; Wang et al., 2023; Cutler et al., 2024; Li et al., 2025) to establish the finite-sample error bounds for a performatively stable mean-field equilibrium.

Assumption 1 (Finite Variance)**.** *There exist constants $\sigma_n \geq 0$ and $\delta \in (0, 1)$ such that, with probability at least $1 - \delta$ over the samples collected during the repeated performative mean field game, for all iterations t and $\bar{m}, m \in \mathcal{M}_{RPMFG}$ and $\theta \in \Theta$:*

$$\left\| \widehat{\mathcal{L}}_\theta(\bar{m}, m) - \mathcal{L}_\theta(\bar{m}, m) \right\|_{H^s(\Theta \times \mathcal{X})} \leq \sigma_n, \tag{11}$$

where $\widehat{\mathcal{L}}_\theta(\bar{m}, m) := \frac{1}{n} \sum_{i=1}^n \ell(x, z_{\theta,i}, \theta, m)$ is the empirical average loss function of a player with θ over n samples and each sample is sampled from the data distribution induced by the type-decision distribution $\bar{m}$, i.e., $z_{\theta,i} \sim \mathcal{D}_\theta(\bar{m})$, and σ_n is the empirical error with n samples at confidence $1 - \delta$.

Based on this assumption, we can now present the following theorem.

Theorem 3. *Suppose Assumption 1 holds, the performative mean field game is α-strongly monotone, the data distribution maps are ε-sensitive for all $\theta \in \Theta$, the loss function $\ell(\cdot)$ is jointly continuous on the product space $\Theta \times \mathcal{X} \times \mathcal{P}(\Theta \times \mathcal{X})$, $\alpha > \varepsilon$ and $s > \frac{d_{\Theta \times \mathcal{X}}}{2}$, then with probability at least $1 - \delta$, the repeated performative mean field game with finite samples produces a sequence of type-decision distributions $\{m^{(0)}, m^{(1)}, m^{(2)}, \cdots, m^{(t)}, \cdots\} \in \mathcal{M}_{RPMFG}$ such that:*

$$\limsup_{t \to \infty} \left\| m^{(t)} - m^{PS} \right\|_{H^{-s}(\Theta \times \mathcal{X})} \leq \frac{\sigma_n}{\alpha - \varepsilon}. \tag{12}$$

The proof of Theorem 3 is given in the Appendix E. Theorem 3 establishes the convergence conditions in the finite-sample setting, with the error bound controlled by α, ε and σ_n. This result provides practical guidance for designing algorithms in the following section.

4 Algorithm

In this section, we propose an adaptive algorithm (Algorithm 1) that achieves a performatively stable mean-field equilibrium without requiring prior knowledge of the data distribution maps $\mathcal{D}_\theta(m)$. Inspired by existing research on mean field games (Zhang et al., 2024; Cui & Koeppl, 2021) and performative prediction (Zhong et al., 2026a) that add regularization terms to construct tractable games, we adopt a similar regularization approach. Specifically, for the original performative mean field game that is ψ-strongly monotone with $\psi \geq 0$, we adaptively add a γ-strongly monotone congestion term $\gamma \cdot C(x, \theta, m)$ to the loss function. At each iteration, $\gamma = \max(0, b\varepsilon - \psi)$ is updated according to the measured ε-sensitivity with $b > 1$, ensuring that the regularized game is α-strongly monotone with $\alpha = \psi + \gamma \geq b\varepsilon > \varepsilon$. By Theorem 3, this regularization ensures convergence to a performatively stable mean-field equilibrium of the regularized game, with error bound $\frac{\sigma_n}{\alpha - \varepsilon}$. When $\psi \geq b\varepsilon$ (so that $\gamma = 0$), this equilibrium coincides with that of the original game. Detailed descriptions are provided in Appendix F.

For the congestion term, we set $C(x, \theta, m) = x \cdot \int_{\mathcal{X}} x m(dx|\theta)$, where $m(dx|\theta)$ is the decision distribution of players with type θ. According to Definition 4, this term is 1-strongly monotone on $\mathcal{M}_{RPMFG}$. For example, when $m(\cdot|\theta)$ is a Dirac measure $\delta_{x^*(\theta)}$, we have the congestion term $C(x, \theta, m) = x \cdot x^*(\theta)$.

Algorithm 1 Performatively Stable Mean-Field Equilibrium

Input: Initial decision distribution $m^{(0)} \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$, 1-strongly monotone congestion term $C(x, \theta, m)$, convergence threshold ρ, initial sensitivity $\varepsilon \geq 0$, initial time step $t = 0$, initial dataset $Z_\theta^{(0)} \sim \mathcal{D}_\theta(m^{(0)})$, Sobolev space dimension $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, parameter $b > 1$,ψ-strongly monotone mean field game $\mathcal{L}_\theta(\bar{m}, m) := \mathbb{E}_{Z_\theta \sim \mathcal{D}_\theta(\bar{m})} \ell(x, Z_\theta, \theta, m)$ with $\psi \geq 0$
Output: The type-decision distribution at performatively stable mean-field equilibrium m^{PS}
repeat
  $t \leftarrow t + 1$;
  $\gamma \leftarrow \max(0, b\varepsilon - \psi)$;
  $\bar{\ell}(x, Z_\theta, \theta, m) \leftarrow \ell(x, Z_\theta, \theta, m) + \gamma \cdot C(x, \theta, m)$;
  $m^{(t)} \in MFE(m^{(t-1)})$ with loss $\bar{\ell}(x, Z_\theta, \theta, m)$;
  Collect data $Z_\theta^{(t)} \sim \mathcal{D}_\theta(m^{(t)})$;
  $\varepsilon \leftarrow \max\left(\varepsilon, \frac{\left\|\mathcal{L}_\theta(m^{(t)}, m^{(t)}) - \mathcal{L}_\theta(m^{(t-1)}, m^{(t)})\right\|_{H^s(\Theta\times\mathcal{X})}}{\left\|m^{(t)} - m^{(t-1)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}}\right)$;
until $\left\|m^{(t)} - m^{(t-1)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \rho$
$m^{PS} \leftarrow m^{(t)}$;

5 Experiment

In this experiment, we evaluate our algorithm on performative mean field games in three scenarios, including quantitative trading, recommendation system, and ride share market[1]. To validate our theoretical results and the algorithm, we set four values for the parameter b in Algorithm 1, where $b = 0, 0.1$ violate the convergence conditions $\alpha > \varepsilon$ in Theorem 3, and $b = 1.1, 10$ satisfy the convergence conditions.

In addition, we compare against three baselines: (1) Static Equilibrium, where players optimize decisions on the initial data distribution without updating; (2) Fixed Regularization $\gamma = 1$, where a constant congestion coefficient $\gamma = 1$ is used; and (3) Heuristic Strategy, where players fit linear data distribution maps (a heuristic rule from (Narang et al., 2023; Tsoy et al., 2025)) using historical decision-data pairs and substitute the fitted maps into the loss function to compute the equilibrium.

5.1 Quantitative Trading

In quantitative trading, $n = 5$ types of traders aim to make decisions close to their market signals z_θ while avoiding congestion from the decisions of other traders. Each type is characterized by a baseline market signal $\theta_i \sim \mathcal{U}[0, 1]^d$ with $d = 5$. Specifically, each trader minimizes $\min_{x\in\mathcal{X}} \|z_\theta - x\|^2 + \eta x^\top \bar{x} - \tau \int_\mathcal{X} m(x|\theta) \log m(x|\theta) dx$, where $x \in \mathcal{X} = [-1, 1]^d$ is the decision of a trader, z_θ is a market signal sample drawn from $\mathcal{D}_\theta(m)$, $\bar{x}$ is the average decision vector of all traders, $\eta x^\top \bar{x}$ is the congestion cost, and $-\tau \int_\mathcal{X} m(x|\theta) \log m(x|\theta) dx$ is the entropy regularization. The data distribution map $\mathcal{D}_\theta(m)$ implies that the collective trading decision distribution induces a nonlinear signal transformation:

$$z_\theta \sim \mathcal{D}_\theta(m) = \theta \odot (\mathbf{1} - (3 + \kappa) \cdot \bar{x} \odot |\bar{x}|) + w, \tag{13}$$

where $\odot$ denotes element-wise multiplication, $\mathbf{1} \in \mathbb{R}^d$ is the all-ones vector, $|\bar{x}|$ is the element-wise absolute value, $w \sim \mathcal{N}(0, \sigma^2 I_d)$, and $(3 + \kappa)$ is the performative effect strength. We set $\kappa \in \{0.25, 0.5, 0.75, 1.0\}$. The implementation details and the calculation of monotonicity, equilibrium, ε-sensitivity can be found in Appendix G.

5.2 Recommendation System

In a recommendation system game, each online merchant with cost θ chooses a price adjustment x to maximize profit. We use the Taobao click dataset[2]. The type distribution is $\theta \sim \mathcal{U}[0.3p, 0.7p]$,

[1]The source code is available at: `https://anonymous.4open.science/r/PerfMFG`.
[2]The dataset is available at: `https://tianchi.aliyun.com/dataset/56`.

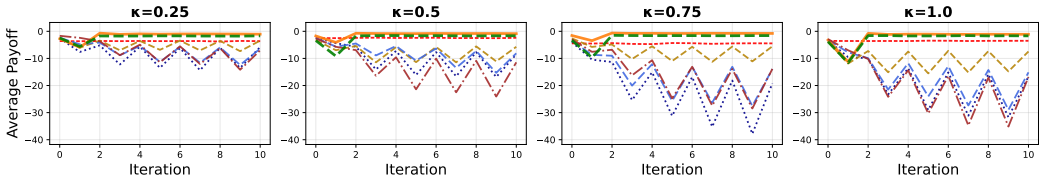

(a) Average payoff in quantitative trading.

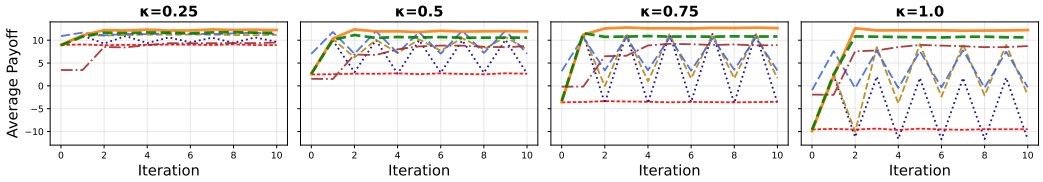

(b) Average payoff in recommendation system.

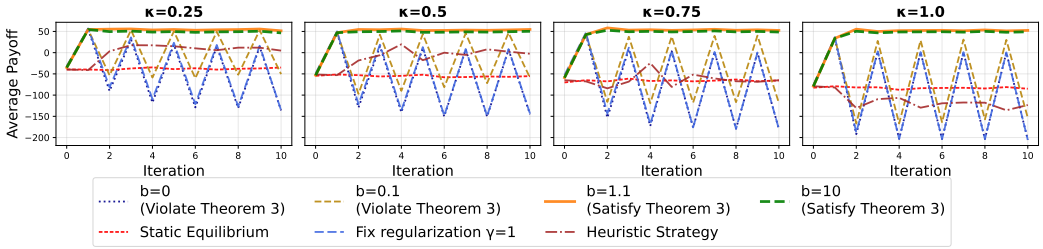

(c) Average payoff in ride share market.

Figure 1: Experimental results with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials. Results with standard deviation are provided in Figs. 9–11 of Appendix J.3.

where p is the average product price. Specifically, each merchant optimizes $\min_{x\in\mathcal{X}} -(1+x)(p-\theta)\cdot z+\eta(1+x)(1+\bar{x})+\frac{\beta}{2}(1+x)^2$, where $x \in \mathcal{X} = [-1, 1]$ is the price adjustment ratio, $z \sim \mathcal{D}(m)$ is the number of clicks, $(1+x)(p-\theta)$ is the profit from a click, $\bar{x}$ is the average price adjustment, $\eta(1+x)(1+\bar{x})$ is the congestion cost, and $\frac{\beta}{2}(1+x)^2$ is the individual adjustment cost. The data distribution map captures the cubic decay in clicks induced by the average price adjustment:

$$z \sim \mathcal{D}(m) = z_0(1-\kappa\bar{x}^3) + w, \tag{14}$$

where z_0 is the baseline click count in the dataset, κ is the performative effect strength, and $w \sim \mathcal{N}(0, (\varsigma z_0)^2)$. We set $\kappa \in \{0.25, 0.5, 0.75, 1.0\}$. The implementation details and the calculation of monotonicity, equilibrium, ε-sensitivity can be found in Appendix H.

5.3 Ride Share Market

In a ride share market, players choose order volume adjustments on Uber and Lyft to maximize utility. We use ride share order data (Narang et al., 2023)[3]. Each player has a type $\theta \sim \text{Beta}(10, 10)$ representing their preference for Uber, with $1-\theta$ the preference for Lyft. Specifically, each player minimizes $\min_{x\in\mathcal{X}} \frac{\zeta}{2}\|x\|^2+p^\top(1+x)-\eta t^\top(1+x)+\beta\mathbf{1}^\top(1+x)\mathbf{1}^\top(1+\bar{x})$, where $x \in \mathcal{X} = [-1, 1]^2$ is the percentage adjustment in order volume on the two platforms, $\frac{\zeta}{2}\|x\|^2$ is the adjustment cost, $p = [p_1, p_2]^\top$ is the price vector induced by average order adjustment $\bar{x}$, p_1 is the price of Uber, p_2 is the price of Lyft, $t = [\theta, 1-\theta]^\top$ is the preference vector, and $\beta\mathbf{1}^\top(1+x)\mathbf{1}^\top(1+\bar{x})$ is the traffic

[3]The dataset is available at:
`https://www.kaggle.com/datasets/brllrb/uber-and-lyft-dataset-boston-ma`.

congestion cost. The data distribution maps reflect that collective order adjustments affect prices:

$$\begin{aligned} p_1 &\sim \mathcal{D}_1(m) = p_{0,1}\big(1 + (1 + 0.4\kappa)\bar{x}_1|\bar{x}_1|\big) + w_1, \\ p_2 &\sim \mathcal{D}_2(m) = p_{0,2}\big(1 + (1 + 0.6\kappa)\bar{x}_2^3\big) + w_2, \end{aligned} \tag{15}$$

where $p_{0,1}$ and $p_{0,2}$ are the baseline prices of two platforms, κ is the performative effect strength, and $w_1, w_2 \sim \mathcal{N}(0, 1)$ are noise terms. We set $\kappa \in \{0.25, 0.5, 0.75, 1.0\}$. The implementation details and the calculation of monotonicity, equilibrium, ε-sensitivity can be found in Appendix I.

5.4 Results and Analysis

We report average payoff, defined as the negative original loss (higher is better), to facilitate comparison across methods in three games. Fig. 1 presents the average payoff across iterations for our algorithm with $b = 0, 0.1, 1.1, 10$ and the three baselines over 10 trials in quantitative trading, the recommendation system, and the ride share market. Additional results are provided in Figs. 2-11 and Table 1 in Appendix J. The results validate the convergence guarantees in Theorem 3 and demonstrate the practical applicability of Algorithm 1.

Across all three games and all κ values, our algorithm with $b = 1.1$ and $b = 10$ (satisfying $\alpha > \varepsilon$) consistently converges to a performatively stable mean-field equilibrium within 3 iterations, where the $\|m^{(t)} - m^{(t-1)}\|_{H^{-s}(\Theta\times\mathcal{X})}$ decreases toward zero as shown in Figs. 2-4 of Appendix J.1. In contrast, $b = 0$ and $b = 0.1$ result in $\alpha < \varepsilon$ in Algorithm 1, which violates the convergence conditions in Theorem 3 and leads to fluctuating payoffs. When $b = 0$, players optimize decisions without regularization. For instance, in the ride share market (Figs. 7 and 8 of Appendix J.2), $b = 1.1$ and $b = 10$ increase order volumes slightly, ensuring the maximization of players' utility after price shifts induced by the type-decision distributions. In contrast, $b = 0$ and $b = 0.1$ oscillate substantially, since $b = 0$ optimizes decisions based on the current price and the congestion term of $b = 0.1$ fails to satisfy $\alpha > \varepsilon$, making their decisions suboptimal after price changes.

As κ increases, the stronger performative effect leads to more severe fluctuations in payoffs for $b = 0$ and $b = 0.1$ at low levels, while our algorithm with $b = 1.1$ and $b = 10$ consistently converges with higher payoffs. Furthermore, although both $b = 1.1$ and $b = 10$ converge, $b = 1.1$ achieves higher payoffs, since a larger b increases the congestion coefficient γ, enforcing stronger regularization and shifting the equilibrium further from the original game.

Static Equilibrium consistently underperforms across all settings, since deploying the initial decision distribution leads to shifts in the data distributions, rendering it suboptimal. Fixed Regularization $\gamma = 1$ performs well when κ is small (e.g., $\kappa = 0.25$ in the recommendation system, Fig. 1 (b)) but its payoffs oscillate as κ increases, since a fixed γ is insufficient to counter strong performative effects. Because the linear heuristic rule in Heuristic Strategy cannot capture the nonlinear performative effects, it optimizes decisions based on a misspecified game structure, leading to relatively poor performance. These comparisons highlight the effectiveness of our algorithm, which ensures convergence to a performatively stable mean-field equilibrium without requiring prior knowledge of the data distribution maps.

6 Conclusion

In this paper, we investigate performative prediction in large-population settings where individual decisions have negligible impact, yet the collective decision distribution meaningfully shapes the data distributions of players. We formalize this setting as a performative mean field game. Leveraging the proposed ε-sensitivity measure, we establish theoretical guarantees for convergence to a performatively stable mean-field equilibrium in finite-sample settings. Accordingly, we design an algorithm that ensures convergence to such an equilibrium. Experiments on three games validate the theoretical results and demonstrate the effectiveness of our algorithm.

A potential limitation of our framework is the setting of full rationality, which can be unrealistic in large-population games. Under bounded rationality, players may not make optimal decisions, leading to a decision distribution that differs from the ideal mean-field equilibrium. To extend our framework to the bounded rationality setting, we will consider that the players select the optimal decision with a probability (Aumann, 1997), where the population decision distribution is a mixture of rational best responses and irrational decisions.

AI USE STATEMENT

In this work, we used generative AI tools for implementing methods. We have not used generative AI tools for developing theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, assisting in the writing of proofs, proposing or refining hypotheses, designing or providing feedback on research methodology or experiments, assisting with the translation of the manuscript, supporting qualitative and thematic data analysis, or interpreting results. Generating synthetic datasets and cleaning or reformatting datasets are not applicable to this work. We have reviewed all AI-assisted work. The AI-assisted code was verified and tested by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

The source code for all experiments is available via the anonymous link provided in Section 5. The complete proofs of all theoretical results, together with the auxiliary definitions and theorems used in the analysis, are provided in the appendix. The detailed experimental settings, including parameter settings, convergence thresholds, and the full game setups for quantitative trading, the recommendation system, and the ride share market, are described in Appendices G, H, and I, with additional results and the statistical significance analysis in Appendix J. The two public datasets used in the experiments, the Taobao click dataset and the Uber and Lyft Boston ride share dataset, are referenced in the footnotes of Section 5.

REFERENCES

Robert J Aumann. Markets with a continuum of traders. *Econometrica*, 32:39–50, 1964.

Robert J Aumann. Rationality and bounded rationality. *Games and Economic Behavior*, 21(1-2): 2–14, 1997.

Claude Berge. *Topological Spaces*. Oliver & Boyd, 1963.

René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications I–II*. Springer, 2018.

Joan Cerdà. *Linear functional analysis*, volume 116. American Mathematical Soc., 2010.

Kai Cui and Heinz Koeppl. Approximately solving mean field games via entropy-regularized deep reinforcement learning. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS'21)*, volume 130, pp. 1909–1917, 13–15 Apr 2021.

Kai Cui, Gökçe Dayanıklı, Mathieu Laurière, Matthieu Geist, Olivier Pietquin, and Heinz Koeppl. Learning discrete-time major-minor mean field games. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI'24)*, volume 38, pp. 9613–9621, 2024.

Joshua Cutler, Mateo Díaz, and Dmitriy Drusvyatskiy. Stochastic approximation with decision-dependent distributions: Asymptotic normality and optimality. *Journal of Machine Learning Research*, 25(90):1–49, 2024.

Szymon Dolecki and Frédéric Mynard. *Convergence foundations of topology*. World Scientific Publishing Company, 2016.

Alexandre Ern and Jean-Luc Guermond. *Finite elements*, volume 1084. Springer, 2021.

Lawrence C Evans and Ronald F Gariepy. *Measure Theory and Fine Properties of Functions*. Chapman and Hall/CRC, 2015. doi: 10.1201/b18333.

Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 1999.

Jöran Friberg. A remarkable collection of babylonian mathematical texts. *Notices of the American Mathematical Society*, 55:1076–1086, 2008.

Diogo A Gomes and Hiroyoshi Mitake. Existence for stationary mean-field games with congestion and quadratic hamiltonians. *Nonlinear Differential Equations and Applications NoDEA*, 22(6): 1897–1910, 2015.

Brandon Gower-Winter, Georg Krempl, Sergey Dragomiretskiy, Tineke Jelsma, and Arno Siebes. Identifying predictions that influence the future: detecting performative concept drift in data streams. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI'25)*, 2025.

Xin Guo, Anran Hu, Renyuan Xu, and Junzi Zhang. Learning mean-field games. In *Advances in Neural Information Processing Systems (NeurIPS 32)*, volume 32, pp. 4967–4977, 2019.

Kun Jin, Tongxin Yin, Zhongzhu Chen, Zeyu Sun, Xueru Zhang, Yang Liu, and Mingyan Liu. Performative federated learning: A solution to model-dependent and heterogeneous distribution shifts. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI'24)*, volume 38, pp. 12938–12946, 2024.

Kun Jin, Tian Xie, Yang Liu, and Xueru Zhang. Addressing polarization and unfairness in performative prediction. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI'26)*, volume 40, pp. 22408–22416, 2026.

Shizuo Kakutani. A Generalization of Brouwer's Fixed Point Theorem. *Duke Mathematical Journal*, 8(3):457–459, September 1941.

Achim Klenke. *Probability theory: a comprehensive course*. Springer, 2008.

Daniel Lacker. Mean field games and interacting particle systems, 2018.

Serge Lang. *Algebra*. Springer, 2002.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2 (1):229–260, 2007.

Hélène Le Cadre, Mandar Datar, Mathis Guckert, and Eitan Altman. Learning market equilibria preserving statistical privacy using performative prediction. *IEEE Transactions on Automatic Control*, 70(11):7125–7140, 2025. doi: 10.1109/TAC.2025.3566920.

Qiang Li, Michal Yemini, and Hoi-To Wai. Clipped SGD algorithms for performative prediction: Tight bounds for stochastic bias and remedies. In *Proceedings of the 42nd International Conference on Machine Learning (ICML'25)*, 2025.

John P Miller, Juan C Perdomo, and Tijana Zrnic. Outside the echo chamber: Optimizing the performative risk. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*, volume 139, pp. 7710–7720, 2021.

Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. Multiplayer performative prediction: Learning in decision-dependent games. *Journal of Machine Learning Research*, 24(202):1–56, 2023.

Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*, volume 119, pp. 7599–7609. PMLR, 2020.

Vasilis Pollatos, Debmalya Mandal, and Goran Radanovic. On corruption-robustness in performative reinforcement learning. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI'25)*, 2025. doi: 10.1609/aaai.v39i19.34196. URL `https://doi.org/10.1609/aaai.v39i19.34196`.

Liangxin Qian, Chang Liu, Xuanyu Cao, Jun Zhao, and Kwok-Yan Lam. Federated bilevel performative prediction. In *Proceedings of the 43rd International Conference on Machine Learning (ICML'26)*, 2026. URL `https://openreview.net/forum?id=7XUocfk7eQ`.

Mitas Ray, Lillian J Ratliff, Dmitriy Drusvyatskiy, and Maryam Fazel. Decision-dependent risk minimization in geometrically decaying dynamic environments. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI'22)*, volume 36, pp. 8081–8088, 2022.

W Rudin. *Principles of mathematical analysis*. McGraw-Hill, United States, 3rd edition, 1976. Includes Index.

Nikita Tsoy, Ivan Kirev, Negin Rahimiyazdi, and Nikola Konstantinov. On the impact of performative risk minimization for binary random variables. In *Proceedings of the 42nd International Conference on Machine Learning (ICML'25)*, 2025.

Xiaolu Wang, Chung-Yiu Yau, and Hoi-To Wai. Network effects in performative prediction games. In *Proceedings of the 40th International Conference on Machine Learning (ICML'23)*, volume 202, pp. 36514–36540, 2023.

Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6):80–83, 1945.

Qiaomin Xie, Zhuoran Yang, Zhaoran Wang, and Andreea Minca. Learning while playing in mean-field games: Convergence and optimality. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*, volume 139, pp. 11436–11447. PMLR, 2021.

Wenjing Yan and Xuanyu Cao. Decentralized noncooperative games with coupled decision-dependent distributions. In *Advances in Neural Information Processing Systems (NeurIPS 37)*, volume 37, pp. 25592–25627, 2024.

Chenyu Zhang, Xu Chen, and Xuan Di. Stochastic semi-gradient descent for learning mean field games with population-aware function approximation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL `https://openreview.net/forum?id=tfO07iz0b9`.

Fengzhuo Zhang, Vincent Y. F. Tan, Zhaoran Wang, and Zhuoran Yang. Learning regularized graphon mean-field games with unknown graphons. *Journal of Machine Learning Research*, 25(372):1–95, 2024.

Guangzheng Zhong, Yang Liu, and Jiming Liu. Towards performatively stable equilibria in decision-dependent games for arbitrary data distribution maps. *Machine Learning*, 115(8):197, 2026a.

Guangzheng Zhong, Yang Liu, Ruichen Liu, and Jiming Liu. Performative prediction in the wild: Adapting to arbitrary data distribution maps. *Machine Learning*, 115(3):42, 2026b.

A Auxiliary Definitions and Theorems

Theorem 4 (Berge's Maximum Theorem (Berge, 1963))**.** *Let $f : \mathcal{X} \times \Theta \to \mathbb{R}$ be a continuous function on the product space $\mathcal{X} \times \Theta$, and $C : \Theta \rightrightarrows \mathcal{X}$ be a compact-valued correspondence. Define the value function $f^* : \Theta \to \mathbb{R}$:*

$$f^*(\theta) = \max\{f(x, \theta) : x \in C(\theta)\}, \tag{16}$$

and the maximizer function $x^* : \Theta \rightrightarrows \mathcal{X}$:

$$x^*(\theta) = \arg\max\{f(x, \theta) : x \in C(\theta)\} = \{x \in C(\theta) : f(x, \theta) = f^*(\theta)\}. \tag{17}$$

If $C(\theta)$ is continuous at θ, then the value function $f^*(\theta)$ is continuous.

Theorem 5 (Dominated Convergence Theorem (Evans & Gariepy, 2015))**.** *Let (Θ, Σ, μ) be a measure space and $\{f_n\}$ be a sequence of measurable functions converging pointwise to f on Θ. If there exists a constant $M > 0$ such that $|f_n(\theta)| \leq M$ for all n and $\theta \in \Theta$, then*

$$\lim_{n\to\infty} \int_\Theta f_n d\mu = \int_\Theta f d\mu. \tag{18}$$

Theorem 6 (Kakutani's Fixed-Point Theorem (Kakutani, 1941))**.** *Suppose K is a convex compact subset of a locally convex topological vector space. Suppose $\Gamma : K \rightrightarrows K$ is a set-valued function satisfying the following conditions:*

1. *$\Gamma(m)$ is nonempty and convex for every $m \in K$.*

2. *The graph* $Gr(\Gamma) = \{(m, m') \in K \times K : m' \in \Gamma(m)\}$ *is closed.*

Then there exists a fixed point, i.e., a point $m \in K$ *such that* $m \in \Gamma(m)$.

Theorem 7 (Riesz-Markov-Kakutani Representation Theorem (Folland, 1999))**.** *Let* A *be a compact space, and for* $\mu \in M(A)$ *and* $f \in C(A)$ *let* $I_\mu(f) = \int f d\mu$*, where* μ *is a measure,* $M(A)$ *is the space of finite signed measures,* f *is a real-valued function, and* $C(A)$ *is the set of continuous functions on* A*. Then the map* $\mu \to I_\mu$ *is an isometric isomorphism from* $M(A)$ *to* $C(A)^*$*, where* $C(A)^*$ *is the topological dual space of* $C(A)$.

Theorem 8 (Alaoglu's Theorem (Folland, 1999))**.** *If* $C(A)$ *is a normed vector space, the closed unit ball* $B^* = \{f \in C(A)^* : \|f\| \leq 1\}$ *in* $C(A)^*$ *is compact in the weak** *topology.*

Theorem 9 (Property of Closed Set (Dolecki & Mynard, 2016))**.** *Any intersection of any family of closed sets is closed (this includes intersections of infinitely many closed sets).*

Theorem 10 (Weierstrass Extreme Value Theorem (Rudin, 1976))**.** *If* K *is a nonempty compact set and* $f : K \to \mathbb{R}$ *is a continuous function, then* f *is bounded and there exist* $p, q \in K$ *such that* $f(p) = \sup_{x \in K} f(x)$ *and* $f(q) = \inf_{x \in K} f(x)$.

Theorem 11 (Heine-Cantor Theorem (Rudin, 1976))**.** *Let* f *be a continuous mapping of a compact metric space* X *into a metric space* Y*. Then* f *is uniformly continuous on* X.

Definition 5 (Weak Convergence of Probability Measures (Klenke, 2008))**.** *Let* A *be a compact space and* $C(A)$ *be the space of continuous bounded functions. A sequence of probability measures* $\{\mu_n\} \subset \mathcal{P}(A)$ *is said to converge weakly to* $\mu \in \mathcal{P}(A)$ *(denoted* $\mu_n \xrightarrow{w} \mu$*) if for every continuous bounded function* $f \in C(A)$*,*

$$\int_A f(a)\mu_n(da) \to \int_A f(a)\mu(da). \tag{19}$$

Definition 6 (Closed Set (Dolecki & Mynard, 2016))**.** *Given a topological space* X*, the following statements are equivalent:*

- *A set* $K \subseteq X$ *is closed in* X.
- K *contains all of its limit points.*

Theorem 12 (Sobolev Embedding Theorem (Ern & Guermond, 2021))**.** *Let* $d \geq 2, s > \frac{d_D}{2}$*,* D *be a bounded Lipschitz domain in* $\mathbb{R}^{d_D}$*. Then the embedding of* $H^s(D)$ *into* $C^0(\bar{D})$ *is continuous, i.e., there is* c *such that* $\|w\|_{C^0} \leq c\|w\|_{H^s(D)}$ *for all* $w \in H^s(D)$*, where* $\bar{D}$ *is the closure of* D*, and* $C^0(\bar{D})$ *consists of continuous functions on* $\bar{D}$.

Definition 7 (Negative-Order Sobolev Space (Ern & Guermond, 2021))**.** *Let* $s > \frac{d_D}{2}$ *and* $D \subset \mathbb{R}^{d_D}$ *be a bounded Lipschitz domain. The space* $H^{-s}(D)$ *is the dual space of* $H^s(D)$ *and is equipped with the norm:*

$$\|T\|_{H^{-s}(D)} := \sup_{w \in H^s(D)} \frac{|\langle w, T\rangle|}{\|w\|_{H^s(D)}}, \tag{20}$$

where T *is a distribution in* D*,* w *is a function in* $H^s(D)$*,* $\|\cdot\|_{H^s(D)}$ *is the norm in space* $H^s(D)$*, and* $\langle w, T\rangle$ *denotes the duality pairing between* $H^s(D)$ *and its dual* $H^{-s}(D)$*. When* T *is a finite signed measure on* D*, the Sobolev Embedding Theorem (Theorem 12) guarantees the map* $w \mapsto \int_D w dT$ *is a bounded linear functional on* $H^s(D)$*. By the definition of the dual norm, we have:*

$$|\langle w, T\rangle| = \left|\int_D w dT\right| \leq \|w\|_{H^s(D)}\|T\|_{H^{-s}(D)}. \tag{21}$$

B Lemma 1 and Proof of Lemma 1

Lemma 1. *If the performative mean field game is* α*-strongly monotone and the loss function* $\ell(\cdot)$ *is jointly continuous on the product space* $\Theta \times \mathcal{X} \times \mathcal{P}(\Theta \times \mathcal{X})$*, then the mean-field equilibrium in each iteration of a repeated performative mean field game exists and is unique. Moreover, let* m^* *be the unique mean-field equilibrium of the game induced by a type-decision distribution* $\bar{m} \in \mathcal{M}_{RPMFG}$*, i.e.,* $m^* \in MFE(\bar{m})$*. Then for all* $m' \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$*:*

$$\int_{\Theta \times \mathcal{X}} \mathcal{L}_\theta(\bar{m}, m^*)(m' - m^*)(d\theta, dx) \geq 0. \tag{22}$$

Proof.

In this proof, to simplify notation, we denote $\mathbb{E}_{Z_\theta \sim \mathcal{D}_\theta(\bar{m})} \ell(x, Z_\theta, \theta, m)$ by $\ell(x, \theta, m)$. Throughout this proof, $\bar{m}$ is an arbitrary and fixed type-decision distribution that induces the data distributions. We prove the existence and uniqueness of the mean-field equilibrium of the game induced by this fixed $\bar{m}$.

We first show the existence of a mean-field equilibrium in each iteration of a repeated performative mean field game.

Let $\mathcal{P}_\mu(\Theta \times \mathcal{X}) := \{m \in \mathcal{P}(\Theta \times \mathcal{X}) : m(d\theta, dx) = \mu(d\theta)\, m(dx|\theta)\}$ denote the set of probability measures on $\Theta \times \mathcal{X}$ that decompose via the fixed type distribution μ.

We consider a map $\Gamma : \mathcal{P}_\mu(\Theta \times \mathcal{X}) \rightrightarrows \mathcal{P}_\mu(\Theta \times \mathcal{X})$ as the best response decision distribution m given the current decision distribution m', which is defined as:

$$\Gamma(m') = \left\{ m \in \mathcal{P}_\mu(\Theta \times \mathcal{X}) : m\left(\left\{ (\theta, x) \in \Theta \times \mathcal{X} : \ell(x, \theta, m') = \min_{x' \in \mathcal{X}} \ell(x', \theta, m') \right\} \right) = 1 \right\}. \tag{23}$$

Since every $m \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$ has marginal distribution μ on Θ, the expected loss satisfies $\int_{\Theta \times \mathcal{X}} \ell(x, \theta, m') m(d\theta, dx) \geq \int_\Theta \min_{x \in \mathcal{X}} \ell(x, \theta, m') \mu(d\theta)$, with equality if and only if m is supported on the minimizers of $\ell(\cdot, \theta, m')$ for μ-almost every θ. Hence $\Gamma(m')$ can be equivalently written as:

$$\Gamma(m') = \left\{ m \in \mathcal{P}_\mu(\Theta \times \mathcal{X}) : \int_{\Theta \times \mathcal{X}} \ell(x, \theta, m') m(d\theta, dx) = \int_\Theta \min_{x \in \mathcal{X}} \ell(x, \theta, m') \mu(d\theta) \right\}. \tag{24}$$

Define $F(m') := \int_\Theta \min_{x \in \mathcal{X}} \ell(x, \theta, m') \mu(d\theta)$. Since ℓ is jointly continuous and $\mathcal{X}$ is compact, then by Berge's Maximum Theorem (Theorem 4), the value function $f^*(\theta, m') := \min_{x \in \mathcal{X}} \ell(x, \theta, m')$ is continuous in (θ, m'). Hence f^* is bounded on the compact domain. By the Dominated Convergence Theorem (Theorem 5), $F(m'_n) \to F(m')$ when $m'_n \to m'$, so F is continuous in m'.

From the definition of mean-field equilibrium (Eq. (4)), m^* is a mean-field equilibrium if and only if it is a fixed point of Γ, i.e., $m^* \in \Gamma(m^*)$.

We now check the conditions of Kakutani's Fixed-Point Theorem (Theorem 6). By the Riesz-Markov-Kakutani Representation Theorem (Theorem 7), the topological dual of $C(\Theta \times \mathcal{X})$ (the space of continuous functions on $\Theta \times \mathcal{X}$), i.e., $C(\Theta \times \mathcal{X})^*$, is isometrically isomorphic to $M(\Theta \times \mathcal{X})$ (the space of signed measures of bounded variation on $\Theta \times \mathcal{X}$). By Alaoglu's Theorem (Theorem 8), the closed unit ball $B^* \subset C(\Theta \times \mathcal{X})^*$ is compact in the weak* topology. Since $\mathcal{P}(\Theta \times \mathcal{X})$ is the space of probability measures, for any $m \in \mathcal{P}(\Theta \times \mathcal{X})$ we have $\|m\| = m(\Theta \times \mathcal{X}) = 1$. Thus $\mathcal{P}(\Theta \times \mathcal{X}) \subset B^* = \{\mu \in C(\Theta \times \mathcal{X})^* : \|\mu\| \leq 1\}$.

Consider $\mathcal{P}(\Theta \times \mathcal{X}) = \{m \in M(\Theta \times \mathcal{X}) : m \geq 0, m(\Theta \times \mathcal{X}) = 1\}$. We now verify that $\{m \geq 0\}$ and $\{m(\Theta \times \mathcal{X}) = 1\}$ are weak*-closed sets:

1. For $m \geq 0$, by the Riesz-Markov-Kakutani Theorem (Theorem 7), m is a nonnegative measure if $\int f dm \geq 0$ for all $f \in C(\Theta \times \mathcal{X})$ with $f \geq 0$. Hence $\{m \geq 0\} = \bigcap_{f \geq 0} \{m : \int f dm \geq 0\}$ is an intersection of weak*-closed half-spaces and therefore is closed in the weak* topology.
2. For $m(\Theta \times \mathcal{X}) = 1$, the function $m \mapsto \int \mathbf{1} dm$ is weak*-continuous because $\mathbf{1} \in C(\Theta \times \mathcal{X})$ (where $\mathbf{1}$ denotes the constant function equal to 1 on $\Theta \times \mathcal{X}$), and $\{m : m(\Theta \times \mathcal{X}) = 1\}$ is the set of all m mapped to 1 by this functional. Hence $\{m(\Theta \times \mathcal{X}) = 1\}$ is closed in the weak* topology.

Thus, by the Property of Closed Set (Theorem 9), the intersection of $\{m \geq 0\}$ and $\{m(\Theta \times \mathcal{X}) = 1\}$, i.e., $\mathcal{P}(\Theta \times \mathcal{X})$, is closed in the weak* topology. Therefore, $\mathcal{P}(\Theta \times \mathcal{X})$ is a closed subset of the compact unit ball B^* and is compact in the weak* topology.

Since the weak* topology is a special case of a polar topology, and polar topologies are locally convex, $C(\Theta \times \mathcal{X})^*$ is a locally convex topological vector space. Thus, $\mathcal{P}(\Theta \times \mathcal{X}) \subset C(\Theta \times \mathcal{X})^*$ is a subset of a locally convex topological vector space.

For any $m_1, m_2 \in \mathcal{P}(\Theta \times \mathcal{X})$ and $a \in [0,1]$, the convex combination $m_a = am_1 + (1-a)m_2$ is a probability measure, which satisfies $m_a(\Theta \times \mathcal{X}) = am_1(\Theta \times \mathcal{X}) + (1-a)m_2(\Theta \times \mathcal{X}) = 1$. Thus, $\mathcal{P}(\Theta \times \mathcal{X})$ is convex.

Combining the above, $\mathcal{P}(\Theta \times \mathcal{X})$ is a compact convex subset of the locally convex topological vector space $C(\Theta \times \mathcal{X})^*$. $\mathcal{P}_\mu(\Theta \times \mathcal{X})$ is convex because for any $m_1, m_2 \in \mathcal{P}_\mu$ and $a \in [0,1]$, the mixture $am_1 + (1-a)m_2$ is a probability measure with $a\mu + (1-a)\mu = \mu$. For closedness, suppose $m_n \in \mathcal{P}_\mu$ converges weakly to m. For any function $f \in C(\Theta)$, we have $\int_{\Theta \times \mathcal{X}} f(\theta) m_n(d\theta, dx) = \int_\Theta f d\mu$. By the definition of weak convergence (Definition 5), we have $\int_{\Theta \times \mathcal{X}} f(\theta) m_n(d\theta, dx) \to \int_{\Theta \times \mathcal{X}} f(\theta) m(d\theta, dx)$. The marginal distribution of m over Θ is also μ, i.e., $m \in \mathcal{P}_\mu$. Hence $\mathcal{P}_\mu(\Theta \times \mathcal{X})$ contains all its limit points and is therefore closed (Definition 6). Thus $\mathcal{P}_\mu(\Theta \times \mathcal{X})$ is a closed convex subset, hence compact and convex. As a result, we can set $K = \mathcal{P}_\mu(\Theta \times \mathcal{X})$ in Theorem 6.

We now check the two properties of the map $\Gamma(m)$:

1. For a fixed $m' \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$, let $S \subset \Theta \times \mathcal{X}$ denote the set of pairs (θ, x) where x is a minimizer of $\ell(x, \theta, m')$, defined as:

$$S = \left\{ (\theta, x) \in \Theta \times \mathcal{X} : \ell(x, \theta, m') = \min_{x' \in \mathcal{X}} \ell(x', \theta, m') \right\}. \tag{25}$$

   $\Gamma(m')$ can be written as $\Gamma(m') = \{m \in \mathcal{P}_\mu(\Theta \times \mathcal{X}) : m(S) = 1\}$. Since $\ell(\cdot)$ is continuous on $\Theta \times \mathcal{X} \times \mathcal{P}(\Theta \times \mathcal{X})$ and the decision space $\mathcal{X}$ is compact, by the Weierstrass Extreme Value Theorem (Theorem 10), $\arg\min_{x \in \mathcal{X}} \ell(x, \theta, m')$ is non-empty for all $\theta \in \Theta$, and S is non-empty. Thus, we can conclude that $\Gamma(m')$ is non-empty.

   For $m_1, m_2 \in \Gamma(m')$ and $a \in [0,1]$, we set $m_a = am_1 + (1-a)m_2$, and we have:

$$m_a(S) = am_1(S) + (1-a)m_2(S) = a \cdot 1 + (1-a) \cdot 1 = 1. \tag{26}$$

   Thus, $m_a \in \Gamma(m')$, which implies that $\Gamma(m')$ is convex.

2. To show that the graph $Gr(\Gamma)$ is closed, suppose $(m'_n, m_n) \subset Gr(\Gamma)$ converges weakly to (m', m), i.e., $(m'_n, m_n) \xrightarrow{w} (m', m)$. Since $m_n \in \Gamma(m'_n)$, by Eq. (24), we have:

$$\int_{\Theta \times \mathcal{X}} \ell(x, \theta, m'_n) m_n(d\theta, dx) = F(m'_n). \tag{27}$$

   Since $\Theta \times \mathcal{X}$ is a compact metric space, $C(\Theta \times \mathcal{X})$ is separable and the weak* topology on the closed unit ball $B^* \subset C(\Theta \times \mathcal{X})^*$ is metrizable. Since $\mathcal{P}(\Theta \times \mathcal{X}) \subset B^*$, its weak* topology is also metrizable. Thus $\Theta \times \mathcal{X} \times \mathcal{P}(\Theta \times \mathcal{X})$ is a compact metric space, and ℓ is uniformly continuous on it by the Heine–Cantor theorem (Theorem 11). Hence $\sup_{\theta,x} |\ell(x, \theta, m'_n) - \ell(x, \theta, m')| \to 0$ as $n \to \infty$. We decompose:

$$\begin{aligned}
&\Big| \int_{\Theta \times \mathcal{X}} \ell(x, \theta, m'_n) m_n(d\theta, dx) - \int_{\Theta \times \mathcal{X}} \ell(x, \theta, m') m(d\theta, dx) \Big| \\
&\leq \int_{\Theta \times \mathcal{X}} \big| \ell(x, \theta, m'_n) - \ell(x, \theta, m') \big| \, m_n(d\theta, dx) + \Big| \int_{\Theta \times \mathcal{X}} \ell(x, \theta, m')(m_n - m)(d\theta, dx) \Big| \\
&\leq \sup_{\theta,x} |\ell(x, \theta, m'_n) - \ell(x, \theta, m')| + \Big| \int_{\Theta \times \mathcal{X}} \ell(x, \theta, m') m_n(d\theta, dx) - \int_{\Theta \times \mathcal{X}} \ell(x, \theta, m') m(d\theta, dx) \Big|.
\end{aligned} \tag{28}$$

   The first term vanishes by uniform continuity. The second term vanishes by the definition of weak convergence (Definition 5) applied to the fixed continuous function $\ell(\cdot, \cdot, m')$. Taking $n \to \infty$ in Eq. (27) yields:

$$\int_{\Theta \times \mathcal{X}} \ell(x, \theta, m') m(d\theta, dx) = F(m'), \tag{29}$$

   which implies $m \in \Gamma(m')$, i.e., $(m', m) \in Gr(\Gamma)$. Hence $Gr(\Gamma)$ contains all its limit points and is therefore closed (Definition 6).

From the above properties, we have that $\Gamma(m)$ is nonempty and convex for every $m \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$, and that the graph $Gr(\Gamma)$ is closed.

Thus, Γ satisfies the conditions of Kakutani's Fixed-Point Theorem. As a result, there exists a fixed point $m^* \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$ such that $m^* \in \Gamma(m^*)$, implying that there exists at least one mean-field equilibrium in each iteration of a repeated performative mean field game.

Now, we show the uniqueness of the mean-field equilibrium in each iteration of a repeated performative mean field game. Suppose $m_1, m_2 \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$ are both mean-field equilibria in an iteration of a repeated performative mean field game, and suppose they are distinct. Then we have:

$$\begin{aligned} \int_{\Theta\times\mathcal{X}} \ell(x,\theta,m_1)m_1(d\theta,dx) - \int_{\Theta\times\mathcal{X}} \ell(x,\theta,m_1)m_2(d\theta,dx) \le 0 \\ \int_{\Theta\times\mathcal{X}} \ell(x,\theta,m_2)m_2(d\theta,dx) - \int_{\Theta\times\mathcal{X}} \ell(x,\theta,m_2)m_1(d\theta,dx) \le 0. \end{aligned} \tag{30}$$

Adding the above two inequalities, we have:

$$\int_{\Theta\times\mathcal{X}} \left(\ell(x,\theta,m_1) - \ell(x,\theta,m_2)\right)(m_1 - m_2)(d\theta,dx) \le 0. \tag{31}$$

Since Eq. (31) contradicts the definition of α-strong monotonicity (Definition 4), there exists at most one mean-field equilibrium in each iteration of a repeated performative mean field game.

Since the mean-field equilibrium m^* is unique, for all $m' \in \mathcal{P}_\mu(\Theta \times \mathcal{X})$, we have:

$$\begin{aligned} &\int_{\Theta\times\mathcal{X}} \ell(x,\theta,m^*)m^*(d\theta,dx) - \int_{\Theta\times\mathcal{X}} \ell(x,\theta,m^*)m'(d\theta,dx) \le 0 \\ \implies &\int_{\Theta\times\mathcal{X}} \ell(x,\theta,m^*)(m' - m^*)(d\theta,dx) \ge 0. \end{aligned} \tag{32}$$

■

C Proof of Theorem 1

Given an α-strongly monotone performative mean field game (Definition 4) and the definition of $MFE(m)$ in Eq. (5), we have, for all $m, m' \in \mathcal{M}_{RPMFG}$, when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$:

$$\begin{aligned} &\alpha \left\| MFE(m) - MFE(m') \right\|^2_{H^{-s}(\Theta\times\mathcal{X})} \\ \le &\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(m, MFE(m)) - \mathcal{L}_\theta(m, MFE(m'))\right)\left(MFE(m) - MFE(m')\right)(d\theta,dx). \end{aligned} \tag{33}$$

By Lemma 1, for all $m, m' \in \mathcal{M}_{RPMFG}$,

$$\begin{aligned} \int_{\Theta\times\mathcal{X}} \mathcal{L}_\theta(m, MFE(m))\left(MFE(m) - MFE(m')\right)(d\theta,dx) \le 0, \\ \int_{\Theta\times\mathcal{X}} \mathcal{L}_\theta(m', MFE(m'))\left(MFE(m) - MFE(m')\right)(d\theta,dx) \ge 0, \end{aligned} \tag{34}$$

and thus we have:

$$\begin{aligned} &\int_{\Theta\times\mathcal{X}} \mathcal{L}_\theta(m, MFE(m))\left(MFE(m) - MFE(m')\right)(d\theta,dx) \\ \le &\int_{\Theta\times\mathcal{X}} \mathcal{L}_\theta(m', MFE(m'))\left(MFE(m) - MFE(m')\right)(d\theta,dx). \end{aligned} \tag{35}$$

Combining Eqs. (33) and (35), when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, we obtain:

$$\begin{aligned} &\alpha \left\| MFE(m) - MFE(m') \right\|^2_{H^{-s}(\Theta\times\mathcal{X})} \\ \le &\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(m, MFE(m)) - \mathcal{L}_\theta(m, MFE(m'))\right)\left(MFE(m) - MFE(m')\right)(d\theta,dx) \\ \le &\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(m', MFE(m')) - \mathcal{L}_\theta(m, MFE(m'))\right)\left(MFE(m) - MFE(m')\right)(d\theta,dx). \end{aligned} \tag{36}$$

By Definition 3, $\mathcal{L}_\theta \in H^s(\Theta \times \mathcal{X})$ with $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, hence by the Sobolev Embedding Theorem (Theorem 12), $\mathcal{L}_\theta$ is also continuous on $\Theta \times \mathcal{X}$. According to the definition of Negative-Order Sobolev Space (Definition 7), a finite measure on $H^{-s}(\Theta \times \mathcal{X})$ is equipped with the norm $\|\cdot\|_{H^{-s}(\Theta\times\mathcal{X})}$. Thus, when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, we have:

$$\begin{aligned}&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(m', MFE(m')) - \mathcal{L}_\theta(m, MFE(m'))\right)(MFE(m) - MFE(m'))\,(d\theta, dx)\\ &\leq \|\mathcal{L}_\theta(m', MFE(m')) - \mathcal{L}_\theta(m, MFE(m'))\|_{H^s(\Theta\times\mathcal{X})} \|MFE(m) - MFE(m')\|_{H^{-s}(\Theta\times\mathcal{X})}.\end{aligned} \tag{37}$$

Combining Eqs. (36) and (37), when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, we obtain:

$$\begin{aligned}&\alpha \|MFE(m) - MFE(m')\|^2_{H^{-s}(\Theta\times\mathcal{X})}\\ &\leq \|\mathcal{L}_\theta(m', MFE(m')) - \mathcal{L}_\theta(m, MFE(m'))\|_{H^s(\Theta\times\mathcal{X})} \|MFE(m) - MFE(m')\|_{H^{-s}(\Theta\times\mathcal{X})}\\ \Longrightarrow &\alpha \|MFE(m) - MFE(m')\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \|\mathcal{L}_\theta(m', MFE(m')) - \mathcal{L}_\theta(m, MFE(m'))\|_{H^s(\Theta\times\mathcal{X})}.\end{aligned} \tag{38}$$

From the definition of the ε-sensitivity (Definition 3), for all $m^*, m_1, m_2 \in \mathcal{M}_{RPMFG}$, when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, we have:

$$\|\mathcal{L}_\theta(m_1, m^*) - \mathcal{L}_\theta(m_2, m^*)\|_{H^s(\Theta\times\mathcal{X})} \leq \varepsilon \|m_1 - m_2\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{39}$$

Since m^*, m_1, m_2 can be any type-decision distributions in the set $\mathcal{M}_{RPMFG}$ as stated in Definition 3, without loss of generality, we set $m^* = MFE(m'), m_1 = m', m_2 = m$. Then, when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, the inequality in Eq. (39) becomes the following:

$$\|\mathcal{L}_\theta(m', MFE(m')) - \mathcal{L}_\theta(m, MFE(m'))\|_{H^s(\Theta\times\mathcal{X})} \leq \varepsilon \|m' - m\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{40}$$

Combining Eqs. (38) and (40), when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, we have:

$$\alpha \|MFE(m) - MFE(m')\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \varepsilon \|m' - m\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{41}$$

Since m, m' can be any type-decision distributions in the set $\mathcal{M}_{RPMFG}$ as stated in Definition 3, without loss of generality, we set $m = m^{(t)}$ and $m' = m^{(t-1)}$, so that $MFE(m) = m^{(t+1)}$, $MFE(m') = m^{(t)}$. Then, when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, we obtain:

$$\left\|m^{(t+1)} - m^{(t)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \frac{\varepsilon}{\alpha} \left\|m^{(t)} - m^{(t-1)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \left(\frac{\varepsilon}{\alpha}\right)^t \left\|m^{(1)} - m^{(0)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{42}$$

■

D Proof of Theorem 2

We first prove that $\{m^{(t)}\}$ is a Cauchy sequence when $\alpha > \varepsilon$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$. For any two positive integers q and p (assume that $q > p$), by repeatedly applying the triangle inequality, we have:

$$\begin{aligned}\left\|m^{(q)} - m^{(p)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} &\leq \left\|m^{(q)} - m^{(p+1)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} + \left\|m^{(p+1)} - m^{(p)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}\\ &\leq \cdots \leq \sum_{r=p}^{q-1} \left\|m^{(r+1)} - m^{(r)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}.\end{aligned} \tag{43}$$

Combining Eqs. (42) and (43), we get:

$$\left\|m^{(q)} - m^{(p)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \sum_{r=p}^{q-1} \left\|m^{(r+1)} - m^{(r)}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \sum_{r=p}^{q-1} \left(\frac{\varepsilon}{\alpha}\right)^r \left\|m^{(1)} - m^{(0)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{44}$$

Let $\mathcal{B} = \frac{\varepsilon}{\alpha}$. Since $\alpha > \varepsilon$, we have $\mathcal{B} < 1$. The sum of the geometric series (Friberg, 2008) is:

$$\sum_{r=p}^{q-1} \mathcal{B}^r = \frac{\mathcal{B}^p \left(1 - \mathcal{B}^{q-p}\right)}{1 - \mathcal{B}} \leq \frac{\mathcal{B}^p}{1 - \mathcal{B}}. \tag{45}$$

Substituting this into Eq. (44), we have:

$$\left\| m^{(q)} - m^{(p)} \right\|_{H^{-s}(\Theta \times \mathcal{X})} \leq \frac{\mathcal{B}^p}{1 - \mathcal{B}} \left\| m^{(1)} - m^{(0)} \right\|_{H^{-s}(\Theta \times \mathcal{X})}. \tag{46}$$

For any $k > 0$, we want to find a number N such that for all $q > p \geq N$, $\left\| m^{(q)} - m^{(p)} \right\|_{H^{-s}(\Theta \times \mathcal{X})} \leq k$. Solving for p, we get:

$$\begin{aligned} &\frac{\mathcal{B}^p}{1 - \mathcal{B}} \left\| m^{(1)} - m^{(0)} \right\|_{H^{-s}(\Theta \times \mathcal{X})} \leq k \iff \mathcal{B}^p \leq \frac{k\left(1 - \mathcal{B}\right)}{\left\| m^{(1)} - m^{(0)} \right\|_{H^{-s}(\Theta \times \mathcal{X})}} \\ \iff & p \log \mathcal{B} \leq \log \left(\frac{k\left(1 - \mathcal{B}\right)}{\left\| m^{(1)} - m^{(0)} \right\|_{H^{-s}(\Theta \times \mathcal{X})}} \right) \iff p \geq \frac{\log \left(\frac{k(1-\mathcal{B})}{\left\| m^{(1)} - m^{(0)} \right\|_{H^{-s}(\Theta \times \mathcal{X})}} \right)}{\log \mathcal{B}}. \end{aligned} \tag{47}$$

Thus, we can choose any positive $N > \frac{\log \left(\frac{k(1-\mathcal{B})}{\left\| m^{(1)} - m^{(0)} \right\|_{H^{-s}(\Theta \times \mathcal{X})}} \right)}{\log \mathcal{B}}$. Since such an N always exists for any $k > 0$, the sequence $\{m^{(t)}\}$ is a Cauchy sequence (Lang, 2002).

Since Eq. (41) holds for any type-decision distributions $m, m' \in \mathcal{M}_{RPMFG}$, the $MFE(\cdot)$ is Lipschitz continuous on $\mathcal{M}_{RPMFG}$. Then, we can have $\lim_{t\to\infty} MFE(m^{(t)}) = MFE(\lim_{t\to\infty} m^{(t)})$.

Using the definition of repeated performative mean field game (Definition 1), i.e., $m^{(t+1)} = MFE(m^{(t)})$, we have: $m^{(\infty)} = MFE(m^{(\infty)})$. Thus, we have $\lim_{t\to\infty} \left\| MFE(m^{(t)}) - m^{(t)} \right\|_{H^{-s}(\Theta \times \mathcal{X})} = 0$. Further, according to the definition of the performatively stable mean-field equilibrium m^{PS} given in Eq. (7), we can conclude that the sequence $\{m^{(t)}\}$ converges to a performatively stable mean-field equilibrium: $\lim_{t\to\infty} m^{(t)} = m^{PS}$. Finally, as $\mathcal{M}_{RPMFG}$ is a closed set that contains all decision distributions during the repeated performative mean field game, it must also contain the limit point $m^{(t\to\infty)}$. Therefore, the performatively stable mean-field equilibrium m^{PS}, being equal to $m^{(t\to\infty)}$, belongs to the set $\mathcal{M}_{RPMFG}$: $m^{PS} \in \mathcal{M}_{RPMFG}$. ■

E Proof of Theorem 3

By the triangle inequality, we have:

$$\begin{aligned} &\left\| \widehat{\mathcal{L}}_\theta \left(m^{(t)}, MFE(m^{(t)}) \right) - \mathcal{L}_\theta \left(m^{PS}, MFE(m^{(t)}) \right) \right\|_{H^s(\Theta \times \mathcal{X})} \\ \leq & \left\| \widehat{\mathcal{L}}_\theta \left(m^{(t)}, MFE(m^{(t)}) \right) - \mathcal{L}_\theta \left(m^{(t)}, MFE(m^{(t)}) \right) \right\|_{H^s(\Theta \times \mathcal{X})} \\ & + \left\| \mathcal{L}_\theta \left(m^{(t)}, MFE(m^{(t)}) \right) - \mathcal{L}_\theta \left(m^{PS}, MFE(m^{(t)}) \right) \right\|_{H^s(\Theta \times \mathcal{X})}, \end{aligned} \tag{48}$$

where $MFE(m^{(t)})$ is the type-decision distribution of mean-field equilibrium under the finite-sample setting, $\widehat{\mathcal{L}}_\theta \left(m^{(t)}, MFE(m^{(t)}) \right)$ represents the average loss at the type-decision distribution $MFE(m^{(t)})$, $\mathcal{L}_\theta \left(m^{(t)}, MFE(m^{(t)}) \right)$ is the expected loss at $MFE(m^{(t)})$ with respect to the true data distributions at time step t, and $\mathcal{L}_\theta \left(m^{PS}, MFE(m^{(t)}) \right)$ is the expected loss at $MFE(m^{(t)})$ with respect to the stable distributions $\mathcal{D}_\theta(m^{PS})$.

From Theorem 2, when $\alpha > \varepsilon$, we know that $m^{PS} \in \mathcal{M}_{RPMFG}$. As stated in Definition 3, when $s > \frac{d_{\Theta \times \mathcal{X}}}{2}$, m, m_1, and m_2 can be any type-decision distribution in the set $\mathcal{M}_{RPMFG}$. Therefore,

when $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, by setting $m = MFE(m^{(t)})$, $m_1 = m^{(t)}$, and $m_2 = m^{PS}$, we have:

$$\left\|\mathcal{L}_\theta\left(m^{(t)}, MFE(m^{(t)})\right) - \mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right)\right\|_{H^s(\Theta\times\mathcal{X})} \leq \varepsilon \left\|m^{(t)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{49}$$

According to Assumption 1, the following inequality holds with probability at least $1-\delta$:

$$\left\|\widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right) - \mathcal{L}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\right\|_{H^s(\Theta\times\mathcal{X})} \leq \sigma_n. \tag{50}$$

Combining Eqs. (48)-(50), when $\alpha > \varepsilon$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, with probability at least $1-\delta$, we have:

$$\left\|\widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right) - \mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right)\right\|_{H^s(\Theta\times\mathcal{X})} \leq \sigma_n + \varepsilon \left\|m^{(t)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{51}$$

Since m^{PS} and $MFE(m^{(t)})$ are the mean-field equilibria under the data distribution induced by m^{PS} and $m^{(t)}$, respectively, and the game is α-strongly monotone, according to Lemma 1, the following inequalities hold. Note that Lemma 1 applies to the empirical loss $\widehat{\mathcal{L}}_\theta$ as well, because $\widehat{\mathcal{L}}_\theta$ uses the same loss function ℓ, the same α-strong monotonicity, and the same continuity properties as $\mathcal{L}_\theta$:

$$\begin{aligned}\int_{\Theta\times\mathcal{X}} \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\left(m^{PS} - MFE(m^{(t)})\right)(d\theta, dx) \geq 0,\\ \int_{\Theta\times\mathcal{X}} \mathcal{L}_\theta\left(m^{PS}, m^{PS}\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx) \geq 0.\end{aligned} \tag{52}$$

Combining the two inequalities in Eq. (52), we have:

$$\begin{aligned}&\int_{\Theta\times\mathcal{X}} \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\left(m^{PS} - MFE(m^{(t)})\right)(d\theta, dx)\\ &+\int_{\Theta\times\mathcal{X}} \mathcal{L}_\theta\left(m^{PS}, m^{PS}\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx)\\ =&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta\left(m^{PS}, m^{PS}\right) - \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx)\\ =&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta\left(m^{PS}, m^{PS}\right) - \mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx)\\ &+\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right) - \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx)\\ \geq&0.\end{aligned} \tag{53}$$

The last inequality in Eq. (53) can be rewritten as follows:

$$\begin{aligned}&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right) - \mathcal{L}_\theta\left(m^{PS}, m^{PS}\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx)\\ \leq&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right) - \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx).\end{aligned} \tag{54}$$

Since the game is α-strongly monotone, we have:

$$\begin{aligned}&\alpha\left\|MFE(m^{(t)}) - m^{PS}\right\|^2_{H^{-s}(\Theta\times\mathcal{X})}\\ \leq&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right) - \mathcal{L}_\theta\left(m^{PS}, m^{PS}\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx).\end{aligned} \tag{55}$$

According to the definition of Negative-Order Sobolev Space (Definition 7), when $\alpha > \varepsilon$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, with probability at least $1-\delta$, we have:

$$\begin{aligned}
&\left|\int_{\Theta\times\mathcal{X}}\left(\mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right) - \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\right)\left(MFE(m^{(t)}) - m^{PS}\right)(d\theta, dx)\right| \\
\leq &\left\|MFE(m^{(t)}) - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})}\left\|\mathcal{L}_\theta\left(m^{PS}, MFE(m^{(t)})\right) - \widehat{\mathcal{L}}_\theta\left(m^{(t)}, MFE(m^{(t)})\right)\right\|_{H^{s}(\Theta\times\mathcal{X})} \\
\leq &\left(\sigma_n + \varepsilon\left\|m^{(t)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})}\right)\left\|MFE(m^{(t)}) - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})},
\end{aligned} \tag{56}$$

where the second inequality follows from Eq. (51).

Combining Eqs. (54)-(56), when $\alpha > \varepsilon$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, with probability at least $1-\delta$, we have:

$$\alpha\left\|MFE(m^{(t)}) - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \sigma_n + \varepsilon\left\|m^{(t)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})}. \tag{57}$$

Since $m^{(t+1)} = MFE(m^{(t)})$, when $\alpha > \varepsilon$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, with probability at least $1-\delta$, we have:

$$\begin{aligned}
\left\|m^{(t+1)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq &\frac{\sigma_n}{\alpha} + \frac{\varepsilon}{\alpha}\left\|m^{(t)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \\
= &\left(\frac{\varepsilon}{\alpha}\right)^t\left\|m^{(1)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})} + \frac{\sigma_n}{\alpha}\sum_{k=0}^{t-1}\left(\frac{\varepsilon}{\alpha}\right)^k.
\end{aligned} \tag{58}$$

The sum of the geometric series (Friberg, 2008) is:

$$\sum_{k=0}^{t-1}\left(\frac{\varepsilon}{\alpha}\right)^k = \frac{1-\left(\frac{\varepsilon}{\alpha}\right)^t}{1-\left(\frac{\varepsilon}{\alpha}\right)}. \tag{59}$$

Since $\alpha > \varepsilon$, we have the following as $t \to \infty$:

$$\begin{aligned}
&\left(\frac{\varepsilon}{\alpha}\right)^t\left\|m^{(1)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \to 0, \\
&\frac{\sigma_n\left(1-\left(\frac{\varepsilon}{\alpha}\right)^t\right)}{\alpha\left(1-\left(\frac{\varepsilon}{\alpha}\right)\right)} \to \frac{\sigma_n}{\alpha-\varepsilon}.
\end{aligned} \tag{60}$$

Thus, when $\alpha > \varepsilon$ and $s > \frac{d_{\Theta\times\mathcal{X}}}{2}$, with probability at least $1-\delta$, we have:

$$\limsup_{t\to\infty}\left\|m^{(t)} - m^{PS}\right\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \frac{\sigma_n}{\alpha-\varepsilon}. \tag{61}$$

■

F Details of Algorithm Description

In each iteration t, the algorithm updates the regularization coefficient γ based on the current estimated ε through $\gamma = \max(0, b\varepsilon - \psi)$ (line 5). Then, for the ψ-strongly monotone game, we add a congestion term $\gamma C(x, \theta, m)$ to the loss function, constructing an α-strongly monotone regularized game with $\alpha = \psi + \gamma$ (line 6). Following this, the algorithm computes the mean-field equilibrium $m^{(t)}$ of the regularized game induced by the previous type-decision distribution $m^{(t-1)}$ (line 7). This step is tractable because the regularized game is strongly monotone, ensuring a unique mean-field equilibrium that can be computed via numerical methods. After deploying the type-decision distribution $m^{(t)}$, the algorithm collects new samples from the data distribution $\mathcal{D}_\theta(m^{(t)})$ induced by $m^{(t)}$ (line 8) and updates the ε estimate (line 9). The procedure repeats until the type-decision distribution stabilizes within a convergence threshold ρ.

Note that line 9 updates ε using consecutive type-decision distributions $m^{(t)}, m^{(t-1)}$, which is sufficient, since the proof of Theorem 1 (Appendix C) applies the ε-sensitivity bound to consecutive type-decision distributions, as shown in Eq. (42).

In our algorithm, the estimate ε is a running maximum, which is monotone non-decreasing and bounded because the ratio in line 9 is bounded by the true ε in Definition 3. Specifically, if the current estimate is smaller than the true ε and the convergence condition is not satisfied, the decisions fluctuate and the estimated ε increases until the condition holds.

When $\gamma > 0$, the algorithm targets the performatively stable equilibrium of the regularized game rather than that of the original game. When the original game does not satisfy the convergence condition $\psi > \varepsilon$, a performatively stable equilibrium may not be achieved during a repeated mean-field game. The regularized game ensures convergence to a performatively stable equilibrium of the regularized game, where the decisions are optimal under the induced data distributions. In addition, since ε is a running maximum, γ is non-decreasing and bounded, and hence converges. As a result, once γ converges, the iteration is that of a fixed α-strongly monotone game with $\alpha = \psi + \gamma \geq b\varepsilon > \varepsilon$, to which Theorem 3 applies, ensuring convergence to a performatively stable equilibrium.

The ε estimate is computed from the collected samples without requiring prior knowledge of the data distribution maps $\mathcal{D}_\theta$. The loss difference $\ell(x, Z_\theta^{(t)}, \theta, m) - \ell(x, Z_\theta^{(t-1)}, \theta, m)$ depends only on the observed data and the loss function ℓ. The H^s norm can be calculated with the loss function ℓ, as detailed in Appendices G-I.

In our experiments, to enable comparison with baselines across iterations, after the stopping criterion $\|m^{(t)} - m^{(t-1)}\|_{H^{-s}(\Theta\times\mathcal{X})} \leq \rho$ is met, we freeze the congestion coefficient γ and the sensitivity estimate ε, but continue to evaluate the equilibrium type-decision distribution $m^{(t)}$ on the induced data distributions and the corresponding payoff at each iteration. This ensures that all methods are compared over the same number of iterations while the algorithm no longer adjusts its regularization parameters after convergence.

G Additional Details of Quantitative Trading

G.1 Implementation Details

In quantitative trading, $n = 5$ types of traders aim to make decisions close to their market signals z_θ while avoiding congestion from the decisions of other traders. The type proportions are drawn from $\mathrm{Dirichlet}(1, \ldots, 1)$, and each type is characterized by a baseline market signal $\theta_i \sim \mathcal{U}[0,1]^d$ with $d = 5$. Specifically, each trader minimizes the following loss:

$$\min_{x\in\mathcal{X}} \|z_\theta - x\|^2 + \eta x^\top \bar{x} - \tau \int_{\mathcal{X}} m(x|\theta) \log m(x|\theta) dx, \tag{62}$$

where $x \in \mathcal{X} = [-1,1]^d$ is the decision of a trader, z_θ is a market signal sample drawn from $\mathcal{D}_\theta(m)$, $\bar{x} = \int_{\Theta\times\mathcal{X}} x m(d\theta, dx)$ is the average decision vector of all traders, $\eta x^\top \bar{x}$ is the congestion cost with $\eta = 0.1$, and $-\tau \int_{\mathcal{X}} m(x|\theta) \log m(x|\theta) dx$ is the entropy regularization with $\tau = 0.2$. The data distribution map $\mathcal{D}_\theta(m)$ implies that the collective trading decision distribution induces a nonlinear signal transformation:

$$z_\theta \sim \mathcal{D}_\theta(m) = \theta \odot (\mathbf{1} - (3+\kappa) \cdot \bar{x} \odot |\bar{x}|) + w, \tag{63}$$

where $\odot$ denotes element-wise multiplication, $\mathbf{1} \in \mathbb{R}^d$ is the all-ones vector, $|\bar{x}|$ is the element-wise absolute value, $w \sim \mathcal{N}(0, \sigma^2 I_d)$ with $\sigma = 0.1$, and $(3+\kappa)$ is the performative effect strength. We set $\kappa \in \{0.25, 0.5, 0.75, 1.0\}$.

The H^{-s} norm is computed via Monte Carlo Fourier integration with $N_\omega = 1000$ frequency samples from a multivariate Cauchy distribution. The convergence threshold is set to $\rho = 10^{-2}$. All experiments are repeated with 10 different random seeds (42–51), and the reported results are averaged over these 10 runs. All experiments were performed on a server running Oracle Linux Server 8.10 with an AMD EPYC 7742 64-Core CPU (256 threads) and 1 TB of RAM, without GPU acceleration. The code is implemented in Python 3.12.10, with key libraries including NumPy 2.5.0, SciPy 1.18.0, and Matplotlib 3.11.0.

G.2 MONOTONICITY

According to the definition, we have:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) - \mathcal{L}_\theta(\bar{m}, m_2)\right)(m_1 - m_2)(d\theta, dx) \\
=&\int_{\Theta\times\mathcal{X}} \left(\left(\|z_\theta - x\|^2 + \eta\cdot x\cdot \bar{x}_1 + \tau H(x)\right) - \left(\|z_\theta - x\|^2 + \eta\cdot x\cdot \bar{x}_2 + \tau H(x)\right)\right)(m_1 - m_2)(d\theta, dx) \\
=&\int_{\Theta\times\mathcal{X}} \eta\cdot x\cdot(\bar{x}_1 - \bar{x}_2)(m_1 - m_2)(d\theta, dx) \\
=&\eta\|\bar{x}_1 - \bar{x}_2\|_2^2 \geq 0,
\end{aligned} \tag{64}$$

The resulting monotonicity is $\psi = 0$.

With entropy regularization $\tau > 0$, the equilibrium conditional distribution for type θ under the quadratic loss is the Gaussian distribution obtained by completing the square:

$$m(x|\theta) \sim \mathcal{N}\left(z_\theta - \frac{\eta}{2}\bar{x}, \frac{\tau}{2}I_d\right), \tag{65}$$

where the type-conditional mean is $\bar{x}(\theta) = \int_{\mathcal{X}} x m(dx|\theta)$.

We consider the congestion term $C(x, \theta, m) = a\cdot x\cdot\int_{\mathcal{X}} x m(dx|\theta)$, where a is a coefficient. Since the conditional decision distribution in mean-field equilibrium is Gaussian with mean $\bar{x}(\theta)$, $\int_{\mathcal{X}} x m(dx|\theta) = \bar{x}(\theta)$ and $C(x, \theta, m) = a\cdot x\cdot\int_{\mathcal{X}} x m(dx|\theta) = a\cdot x\cdot\bar{x}(\theta)$. In this case, the monotonicity of the game becomes:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) + \gamma C(x, \theta, m_1) - \mathcal{L}_\theta(\bar{m}, m_2) - \gamma C(x, \theta, m_2)\right)(m_1 - m_2)(d\theta, dx) \\
=&\eta\|\bar{x}_1 - \bar{x}_2\|_2^2 + \gamma\cdot a\cdot\int_\Theta \|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2\mu(d\theta) \\
\geq&\gamma\cdot a\cdot\int_\Theta \|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2\mu(d\theta).
\end{aligned} \tag{66}$$

The H^{-s} norm of the difference between two equilibrium measures is computed via Fourier transformation (Cerdà, 2010). Specifically, for any function ϕ with dimension D, we have:

$$\|\phi\|_{H^{-s}(\mathbb{R}^D)}^2 = \int_{\mathbb{R}^D}\left(1 + |\xi|^2\right)^{-s}\left|\hat{\phi}(\xi)\right|^2 d\xi. \tag{67}$$

For an equilibrium measure m, with Gaussian conditional distribution $m(\cdot|\theta)$, the Fourier transform is

$$\widehat{m}(\xi, \eta) = \int_\Theta e^{-i\xi\cdot\theta}\left(\int_{\mathcal{X}} e^{-i\eta\cdot x} m(dx|\theta)\right)\mu(d\theta). \tag{68}$$

The Gaussian characteristic function satisfies $\int_{\mathcal{X}} e^{-i\eta\cdot x} m(dx|\theta) = e^{-i\eta\cdot\bar{x}(\theta)}\cdot e^{-\tau\|\eta\|_2^2/4}$. Since $e^{-\tau\|\eta\|_2^2/4} \leq 1$, the Fourier difference satisfies

$$\begin{aligned}
|\widehat{m}_1(\xi, \eta) - \widehat{m}_2(\xi, \eta)| &= e^{-\tau\|\eta\|_2^2/4}\left|\int_\Theta e^{-i\xi\cdot\theta}\left(e^{-i\eta\cdot\bar{x}_1(\theta)} - e^{-i\eta\cdot\bar{x}_2(\theta)}\right)\mu(d\theta)\right| \\
&\leq \left|\int_\Theta e^{-i\xi\cdot\theta}\left(e^{-i\eta\cdot\bar{x}_1(\theta)} - e^{-i\eta\cdot\bar{x}_2(\theta)}\right)\mu(d\theta)\right|.
\end{aligned} \tag{69}$$

Applying Jensen's inequality and the bound $|e^{-it} - 1| = 2|\sin(t/2)| \leq |t|$, we obtain

$$\begin{aligned}
|\widehat{m}_1(\xi, \eta) - \widehat{m}_2(\xi, \eta)|^2 &\leq \int_\Theta \left|e^{-i\eta\cdot\bar{x}_1(\theta)} - e^{-i\eta\cdot\bar{x}_2(\theta)}\right|^2\mu(d\theta) \\
&= \int_\Theta 4\sin^2\left(\frac{\eta\cdot(\bar{x}_1(\theta) - \bar{x}_2(\theta))}{2}\right)\mu(d\theta) \\
&\leq \int_\Theta \|\eta\|_2^2\|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2\mu(d\theta).
\end{aligned} \tag{70}$$

Substituting into $\|m_1 - m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}$, we have:

$$\begin{aligned}\|m_1 - m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})} &= \int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} |\widehat{m}_1(\xi,\eta) - \widehat{m}_2(\xi,\eta)|^2 \, d\xi d\eta \\ &\le \int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} \int_\Theta \|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2 \|\eta\|_2^2 \mu(d\theta) d\xi d\eta \\ &= \int_\Theta \|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2 \mu(d\theta) \int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} \|\eta\|_2^2 d\xi d\eta.\end{aligned} \tag{71}$$

For the calculation of the integral, we transform ξ, η into a spherical coordinate system and integrate via Beta functions. Note that the presence of $\|\eta\|_2^2$ in the integrand requires $s > \frac{d_{\Theta\times\mathcal{X}}}{2} + 1$ for the Gamma function $\Gamma(s - \frac{d_\Theta + d_\mathcal{X}}{2} - 1)$. This stronger condition is satisfied by all experimental s values chosen in this work:

$$\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} \|\eta\|_2^2 d\xi d\eta = \frac{\pi^{\frac{d_\Theta + d_\mathcal{X}}{2}} d_\mathcal{X} \Gamma(s - \frac{d_\Theta + d_\mathcal{X}}{2} - 1)}{2\Gamma(s)}. \tag{72}$$

Since $d_\Theta = d$, $d_\mathcal{X} = d$, $d_{\Theta\times\mathcal{X}} = 2d$, with $d = 5$ and $s = 10$:

$$I(10) = \frac{\pi^5 \cdot 5 \cdot \Gamma(10 - 5 - 1)}{2\Gamma(10)} = \frac{5\pi^5\Gamma(4)}{2\Gamma(10)}. \tag{73}$$

Thus, when $\frac{\pi^{\frac{d_\Theta + d_\mathcal{X}}{2}} d_\mathcal{X} \Gamma(s - \frac{d_\Theta + d_\mathcal{X}}{2} - 1)}{2\Gamma(s)} \le 1$, we have:

$$\int_\Theta \|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2 \mu(d\theta) \ge \|m_1 - m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}. \tag{74}$$

In this experiment, we have the coefficient $a = \frac{5\pi^5\Gamma(4)}{2\Gamma(10)}$:

$$\begin{aligned}&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) + \gamma C(x,\theta,m_1) - \mathcal{L}_\theta(\bar{m}, m_2) - \gamma C(x,\theta,m_2)\right)(m_1 - m_2)(d\theta, dx) \\ \ge&\gamma \cdot a \cdot \int_\Theta \|\bar{x}_1(\theta) - \bar{x}_2(\theta)\|_2^2 \mu(d\theta) \\ \ge&\frac{\gamma a \|m_1 - m_2\|^2_{H^{-10}(\Theta\times\mathcal{X})}}{\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} (1 + \|\xi\|^2 + \|\eta\|^2)^{-10} \|\eta\|^2 d\xi d\eta} \\ =&\gamma \|m_1 - m_2\|^2_{H^{-10}(\Theta\times\mathcal{X})}.\end{aligned} \tag{75}$$

As a result, after adding the congestion term $\gamma C(x,\theta,m)$ with $a = \frac{5\pi^5\Gamma(4)}{2\Gamma(10)}$ and setting $s = 10$, the quantitative trading game becomes γ-strongly monotone.

G.3 Calculation of Mean-Field Equilibrium

With entropy regularization $-\tau H(m(\cdot|\theta))$ where $H(m(\cdot|\theta)) = -\int_\mathcal{X} m(x|\theta) \log m(x|\theta) dx$, the equilibrium conditional distribution for type θ is the Gibbs distribution $m(x|\theta) \propto \exp(\frac{-\ell(x,\theta,m)}{\tau})$. For the quadratic loss, completing the square yields the Gaussian distribution $\mathcal{N}(z_\theta - \frac{\eta}{2}\bar{x}, \frac{\tau}{2} I_d)$. The type-conditional mean $\bar{x}(\theta) = \int_\mathcal{X} x m(dx|\theta)$ therefore satisfies the fixed-point equation:

$$\bar{x}(\theta) = z_\theta - \frac{\eta}{2}\bar{x}. \tag{76}$$

The global mean $\bar{x} = \sum_{i=1}^n p_i \bar{x}(\theta_i)$ with n discrete types (weights p_i, $\sum_{i=1}^n p_i = 1$) satisfies $\bar{x} = \bar{z} - \frac{\eta\bar{x}}{2}$, where $\bar{z} = \sum_{i=1}^n p_i z_{\theta_i}$. Rearranging:

$$\bar{x} = \frac{\bar{z}}{1 + \frac{\eta}{2}}. \tag{77}$$

The per-type optimal decision is $x_i^* = \bar{x}(\theta_i) = z_{\theta_i} - \frac{\eta\bar{x}}{2}$. To respect the bounded decision space $\mathcal{X} = [-R, R]^d$ with $R = 1$, each decision coordinate is clipped and the global mean recomputed: $x_i^* \leftarrow \mathrm{clip}(x_i^*,\ -R,\ R)$, $\bar{x} \leftarrow \sum_{i=1}^n p_i x_i^*$.

G.4 Calculation of ε-Sensitivity

From the loss function, we have:

$$\left\|\mathcal{L}_\theta(m^{(t-1)}, m) - \mathcal{L}_\theta(m^{(t)}, m)\right\|_{H^s(\Theta\times\mathcal{X})} = \left\|\|z_\theta^{(t-1)} - x\|_2^2 - \|z_\theta^{(t)} - x\|_2^2\right\|_{H^s(\Theta\times\mathcal{X})}. \tag{78}$$

We have:

$$\varepsilon = \max\left(\varepsilon, \frac{\left\|\|z_\theta^{(t-1)} - x\|_2^2 - \|z_\theta^{(t)} - x\|_2^2\right\|_{H^s(\Theta\times\mathcal{X})}}{\left\|m^{(t)} - m^{(t-1)}\right\|_{H^{-s}(\Theta\times\mathcal{X})}}\right). \tag{79}$$

We set $f = \|z_\theta^{(t-1)} - x\|_2^2 - \|z_\theta^{(t)} - x\|_2^2 = \Delta(z^2)_\theta - 2x^\top \Delta z_\theta$, where $\Delta(z^2)_\theta = \|z_\theta^{(t-1)}\|_2^2 - \|z_\theta^{(t)}\|_2^2$, and $\Delta z_\theta = z_\theta^{(t-1)} - z_\theta^{(t)}$. According to the definition of the Sobolev norm, we have

$$\|f\|_{H^s(\Theta\times\mathcal{X})} = \left(\sum_{|\beta|\le s} \|D^\beta f\|_{L^2(\Theta\times\mathcal{X})}^2\right)^{\frac{1}{2}}, \tag{80}$$

where $D^\beta f$ is a weak β-th partial derivative of f.

Since f is affine in x, all x-derivatives of order $|\beta| \ge 2$ are identically zero. Moreover, the type space $\Theta = \{\theta_1, \ldots, \theta_n\}$ is discrete, so the integral over Θ in the Sobolev norm reduces to a weighted sum $\sum_i p_i$. Consequently, no θ-derivatives enter the computation, and only the L^2 term and first-order x-derivatives contribute. For $s \ge 2$, we have:

$$\|f\|_{H^s(\Theta\times\mathcal{X})} = \left(\|f\|_{L^2(\Theta\times\mathcal{X})}^2 + \sum_{j=1}^d \|\partial_{x_j} f\|_{L^2(\Theta\times\mathcal{X})}^2\right)^{\frac{1}{2}}. \tag{81}$$

Since $\mathcal{X} = [-R, R]^d$, we have $\int_\mathcal{X} 1dx = (2R)^d$, $\int_\mathcal{X} x_j dx = 0$, $\int_\mathcal{X} x_j^2 dx = \frac{(2R)^d R^2}{3}$, and $\int_\mathcal{X} x_j x_k dx = 0$ for $j \ne k$. For discrete types, the integral over Θ reduces to a weighted sum with weights p_i. For each term, we have:

$$\begin{aligned} \|f\|_{L^2(\Theta\times\mathcal{X})}^2 &= (2R)^d \sum_{i=1}^n p_i \left[(\Delta(z^2)_{\theta_i})^2 + \frac{4R^2}{3}\|\Delta z_{\theta_i}\|_2^2\right], \\ \sum_{j=1}^d \|\partial_{x_j} f\|_{L^2(\Theta\times\mathcal{X})}^2 &= 4(2R)^d \sum_{i=1}^n p_i \|\Delta z_{\theta_i}\|_2^2. \end{aligned} \tag{82}$$

Combining:

$$\|f\|_{H^s(\Theta\times\mathcal{X})}^2 = (2R)^d \sum_{i=1}^n p_i \left[(\Delta(z^2)_{\theta_i})^2 + \left(\frac{4R^2}{3} + 4\right)\|\Delta z_{\theta_i}\|_2^2\right], \tag{83}$$

which is computed in $\mathcal{O}(n)$ without requiring any knowledge of the distribution map.

From the monotonicity calculation, the H^{-s} norm is computed via Monte Carlo Fourier integration with multivariate Cauchy importance sampling (Cerdà, 2010):

$$\|m^{(t)} - m^{(t-1)}\|_{H^{-s}(\Theta\times\mathcal{X})}^2 = \int_{\mathbb{R}^{2d}} (1 + \|\xi\|_2^2 + \|\eta\|_2^2)^{-s} \left|\sum_{i=1}^n p_i e^{-i\xi\cdot\theta_i}\left(e^{-i\eta\cdot x_i^{(t)}} - e^{-i\eta\cdot x_i^{(t-1)}}\right)\right|^2 d\xi d\eta. \tag{84}$$

H Additional Details of Recommendation System

H.1 Implementation Details

In a recommendation system game, each online merchant with cost θ chooses a price adjustment x to maximize profit. We use the Taobao click dataset[4], from which we select the 36 product categories with at least 5,000 clicks, each treated as an independent mean field game instance. The type distribution is $\theta \sim \mathcal{U}[0.3p, 0.7p]$, where p is the average product price of the category. Specifically, each merchant optimizes the following problem:

$$\min_{x\in\mathcal{X}} -(1+x)(p-\theta)\cdot z + \eta(1+x)(1+\bar{x}) + \frac{\beta}{2}(1+x)^2, \tag{85}$$

where $x \in \mathcal{X} = [-1, 1]$ is the price adjustment ratio, $z \sim \mathcal{D}(m)$ is the number of clicks, $(1+x)(p-\theta)$ is the profit from a click, $\bar{x} = \int_{\Theta\times\mathcal{X}} xm(d\theta, dx)$ is the average price adjustment, $\eta(1+x)(1+\bar{x})$ is the congestion cost with $\eta = 0.2$, and $\frac{\beta}{2}(1+x)^2$ is the individual adjustment cost with $\beta = 0.1 \cdot p \cdot z_0$, which is scaled by the category-specific average price and baseline click count to ensure that the adjustment cost is proportional to the profit scale of each category. The data distribution map captures the cubic decay in clicks induced by the average price adjustment:

$$z \sim \mathcal{D}(m) = z_0(1 - \kappa\bar{x}^3) + w, \tag{86}$$

where z_0 is the baseline click count normalized by the average number of clicks in the dataset, κ is the performative effect strength, and $w \sim \mathcal{N}(0, (\varsigma z_0)^2)$ with $\varsigma = 0.01$. We set $\kappa \in \{0.25, 0.5, 0.75, 1.0\}$.

We draw $N_\theta = 500$ type samples from $\mathcal{U}[0.3p, 0.7p]$ and $N_\omega = 500$ frequency samples for the H^{-s} norm computation, following the same Monte Carlo Fourier integration procedure as in Appendix G. The convergence threshold is set to $\rho = 5 \times 10^{-2}$. All experiments are repeated with 10 different random seeds (42–51), and the reported results are averaged over these 10 runs. All experiments were performed on a server running Oracle Linux Server 8.10 with an AMD EPYC 7742 64-Core CPU (256 threads) and 1 TB of RAM, without GPU acceleration. The code is implemented in Python 3.12.10, with key libraries including NumPy 2.5.0, SciPy 1.18.0, and Matplotlib 3.11.0.

H.2 Monotonicity

According to the definition, we have:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) - \mathcal{L}_\theta(\bar{m}, m_2)\right)(m_1 - m_2)(d\theta, dx) \\
=&\int_{\Theta\times\mathcal{X}} \eta(1+x)(\bar{x}_1 - \bar{x}_2)\,(m_1 - m_2)\,(d\theta, dx) \\
=&\eta(\bar{x}_1 - \bar{x}_2)\int_{\Theta\times\mathcal{X}} (1+x)\,(m_1 - m_2)\,(d\theta, dx) \\
=&\eta\|\bar{x}_1 - \bar{x}_2\|_2^2.
\end{aligned} \tag{87}$$

Since the decision distribution given type θ in the mean-field equilibrium is a Dirac measure, every measure $m \in \mathcal{M}_{RPMFG}$ satisfies that, for all measurable sets $A \subset \Theta$ and $B \subset \mathcal{X}$, $m(A \times B) = \int_A \delta_{x^*(\theta)}(B)\mu(d\theta)$, where μ is the fixed probability measure over the type space Θ, and $\delta_{x^*(\theta)}$ is the Dirac measure on $\mathcal{X}$ at the point $x^*(\theta)$ for type θ.

We consider the congestion term $C(x, \theta, m) = a \cdot x \cdot \int_{\mathcal{X}} xm(dx|\theta)$, where a is a coefficient. Since the decision distribution in mean-field equilibrium is a Dirac measure, $\int_{\mathcal{X}} xm(dx|\theta) = x^*(\theta)$ and

[4] Available at `https://tianchi.aliyun.com/dataset/56`.

$C(x,\theta,m) = a\cdot x\cdot\int_{\mathcal{X}} xm(dx|\theta) = a\cdot x\cdot x^*(\theta)$. In this case, the monotonicity of the game becomes:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) + \gamma C(x,\theta,m_1) - \mathcal{L}_\theta(\bar{m}, m_2) - \gamma C(x,\theta,m_2)\right)(m_1-m_2)(d\theta,dx)\\
=&\eta\|\bar{x}_1-\bar{x}_2\|_2^2 + \gamma\cdot a\cdot\int_\Theta \|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)\\
=&\eta\|\bar{x}_1-\bar{x}_2\|_2^2 + \gamma\cdot a\cdot\int_\Theta \left\|\int_{\mathcal{X}} x'\left(\delta_{x_1^*(\theta)}(dx') - \delta_{x_2^*(\theta)}\right)(dx')\right\|_2^2\mu(d\theta).
\end{aligned} \tag{88}$$

Since every measure $m\in\mathcal{M}_{RPMFG}$ satisfies that, for all measurable sets $A\subset\Theta$ and $B\subset\mathcal{X}$, $m(A\times B) = \int_A \delta_{x^*(\theta)}(B)\mu(d\theta)$, $\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}$ can be calculated using Fourier transformation (Cerdà, 2010). Specifically, for any function ϕ with dimension D, we have:

$$\|\phi\|^2_{H^{-s}(\mathbb{R}^D)} = \int_{\mathbb{R}^D}\left(1+|\xi|^2\right)^{-s}\left|\hat{\phi}(\xi)\right|^2 d\xi. \tag{89}$$

Then, we have:

$$\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})} = \int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}|\hat{m}_1(\xi,\eta)-\hat{m}_2(\xi,\eta)|^2\,d\xi d\eta. \tag{90}$$

Since every measure $m\in\mathcal{M}_{RPMFG}$ satisfies that, for all measurable sets $A\subset\Theta$ and $B\subset\mathcal{X}$, $m(A\times B) = \int_A \delta_{x^*(\theta)}(B)\mu(d\theta)$, we have

$$\hat{m}(\xi,\eta) = \int_{\Theta\times\mathcal{X}} e^{-i(\xi\theta+\eta x)}m(d\theta,dx) = \int_\Theta e^{-i(\xi\theta+\eta x^*(\theta))}\mu(d\theta). \tag{91}$$

Then, we have:

$$\begin{aligned}
&\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}\\
=&\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\left|\int_\Theta e^{-i(\xi\theta+\eta x_1^*(\theta))}\mu(d\theta) - \int_\Theta e^{-i(\xi\theta+\eta x_2^*(\theta))}\mu(d\theta)\right|^2 d\xi d\eta\\
=&\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right)\mu(d\theta)\right|^2 d\xi d\eta.
\end{aligned} \tag{92}$$

Now, we focus on $\left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right)\mu(d\theta)\right|^2$. We have:

$$\begin{aligned}
\left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right)\mu(d\theta)\right|^2 &\le \int_\Theta \left|e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right)\right|^2\mu(d\theta)\\
&= \int_\Theta \left|e^{-i\xi\theta}\right|^2\left|e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right|^2\mu(d\theta).
\end{aligned} \tag{93}$$

According to Euler's formula, we have:

$$\left|e^{-i\xi\theta}\right|^2 = |\cos(\xi\theta) - i\sin(\xi\theta)|^2 = \cos(\xi\theta)^2 + (-\sin(\xi\theta))^2 = \cos^2(\xi\theta)+\sin^2(\xi\theta) = 1. \tag{94}$$

Thus, $\left|e^{-i\xi\theta}\right| = 1$.

For $\left|e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right|^2$, we have:

$$\begin{aligned}
&\left|e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)}\right|^2\\
=&(e^{-i\eta x_1^*(\theta)} - e^{-i\eta x_2^*(\theta)})\cdot(e^{i\eta x_1^*(\theta)} - e^{i\eta x_2^*(\theta)})\\
=&e^{-i\eta x_1^*(\theta)+i\eta x_1^*(\theta)} - e^{-i\eta x_1^*(\theta)+i\eta x_2^*(\theta)} - e^{-i\eta x_2^*(\theta)+i\eta x_1^*(\theta)} + e^{-i\eta x_2^*(\theta)+i\eta x_2^*(\theta)}\\
=&1+1-(e^{-i\eta x_1^*(\theta)+i\eta x_2^*(\theta)} + e^{-i\eta x_2^*(\theta)+i\eta x_1^*(\theta)})\\
=&2-\left(e^{i\eta(x_1^*(\theta)-x_2^*(\theta))} + e^{-i\eta(x_1^*(\theta)-x_2^*(\theta))}\right).
\end{aligned} \tag{95}$$

According to Euler's formula, we have:

$$e^{i\eta(x_1^*(\theta)-x_2^*(\theta))}+e^{-i\eta(x_1^*(\theta)-x_2^*(\theta))}=2\cos(\eta(x_1^*(\theta)-x_2^*(\theta))). \tag{96}$$

Then, we have:

$$\left|e^{-i\eta x_1^*(\theta)}-e^{-i\eta x_2^*(\theta)}\right|^2=2-2\cos(\eta(x_1^*(\theta)-x_2^*(\theta))). \tag{97}$$

According to the trigonometric identity $1-\cos\theta=2\sin^2\frac{\theta}{2}$, we have:

$$\left|e^{-i\eta x_1^*(\theta)}-e^{-i\eta x_2^*(\theta)}\right|^2=4\sin^2\frac{\eta(x_1^*(\theta)-x_2^*(\theta))}{2}. \tag{98}$$

Then, we have:

$$\begin{aligned}\left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)}-e^{-i\eta x_2^*(\theta)}\right)\mu(d\theta)\right|^2&\le\int_\Theta\left|e^{-i\xi\theta}\right|^2\left|e^{-i\eta x_1^*(\theta)}-e^{-i\eta x_2^*(\theta)}\right|^2\mu(d\theta)\\&=\int_\Theta 4\sin^2\frac{\eta(x_1^*(\theta)-x_2^*(\theta))}{2}\mu(d\theta).\end{aligned} \tag{99}$$

Since $|\sin x|\le|x|$, we have

$$\sin^2 x\le x^2\Longrightarrow\sin^2\frac{\eta(x_1^*(\theta)-x_2^*(\theta))}{2}\le\left|\frac{\eta(x_1^*(\theta)-x_2^*(\theta))}{2}\right|^2\le\frac{\|\eta\|_2^2\|x_1^*(\theta)-x_2^*(\theta)\|_2^2}{4}. \tag{100}$$

Thus, we have:

$$\begin{aligned}\left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)}-e^{-i\eta x_2^*(\theta)}\right)\mu(d\theta)\right|^2&\le\int_\Theta 4\sin^2\frac{\eta(x_1^*(\theta)-x_2^*(\theta))}{2}\mu(d\theta)\\&\le\int_\Theta\|\eta\|_2^2\|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta).\end{aligned} \tag{101}$$

Substituting into $\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}$, we have:

$$\begin{aligned}&\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}\\&=\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x_1^*(\theta)}-e^{-i\eta x_2^*(\theta)}\right)\mu(d\theta)\right|^2d\xi d\eta\\&\le\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\int_\Theta\|\eta\|_2^2\|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)d\xi d\eta\\&=\int_\Theta\|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2d\xi d\eta.\end{aligned} \tag{102}$$

Thus, when $\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2d\xi d\eta\le 1$, we have:

$$\int_\Theta\|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)\ge\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}. \tag{103}$$

For the calculation of $\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2d\xi d\eta$, we transform ξ,η into a spherical coordinate system:

$$\xi\to(\Xi,\Phi_1,\cdots\Phi_{d_\Theta-1}),\eta\to(\hat{\eta},\Psi_1,\cdots,\Psi_{d_\mathcal{X}-1}). \tag{104}$$

Then, we have:

$$\|\xi\|_2^2=\Xi^2,\|\eta\|_2^2=\hat{\eta}^2. \tag{105}$$

According to integration in polar coordinates in $\mathbb{R}^n$, we have:

$$\int_{\mathbb{R}^n} f(x)dx = \omega_n \int_0^\infty g(r) r^{n-1} dr, \tag{106}$$

where $r = \|x\|_2$, $f(x) = g(\|x\|_2)$, $\omega_n = \frac{2\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2})}$ is the surface area of the unit $(n-1)$-sphere of radius 1 embedded in n-dimensional Euclidean space, $\Gamma(\cdot)$ is the Gamma function.

Using this integration in polar coordinates in $\mathbb{R}^n$, we have

$$\begin{aligned}\int_{\mathbb{R}^{d_\Theta \times \mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} \|\eta\|_2^2 d\xi d\eta =& \omega_{d_\Theta}\omega_{d_\mathcal{X}} \int_0^\infty \int_0^\infty \frac{\hat{\eta}^2}{(1+\hat{\eta}^2+\Xi^2)^s} \hat{\eta}^{d_\mathcal{X}-1} \Xi^{d_\Theta - 1} d\hat{\eta} d\Xi \\ =& \omega_{d_\Theta}\omega_{d_\mathcal{X}} \int_0^\infty \int_0^\infty \frac{\hat{\eta}^{d_\mathcal{X}+1}}{(1+\hat{\eta}^2+\Xi^2)^s} \Xi^{d_\Theta - 1} d\hat{\eta} d\Xi.\end{aligned} \tag{107}$$

Let $u = \Xi^2$, then $\Xi = u^{\frac{1}{2}}, d\Xi = \frac{1}{2}u^{-\frac{1}{2}}du$, and we have:

$$\begin{aligned}\int_0^\infty \frac{1}{(1+\hat{\eta}^2+\Xi^2)^s} \Xi^{d_\Theta - 1} d\Xi =& \int_0^\infty \frac{u^{\frac{d_\Theta - 1}{2}}}{(1+\hat{\eta}^2+u)^s} \frac{1}{2} u^{-\frac{1}{2}} du \\ =& \frac{1}{2}\int_0^\infty \frac{u^{\frac{d_\Theta}{2}-1}}{(1+\hat{\eta}^2+u)^s} du.\end{aligned} \tag{108}$$

Let $t = \frac{u}{1+\hat{\eta}^2}$, then $u = (1+\hat{\eta}^2)t, du = 1+\hat{\eta}^2 dt$, we have:

$$\begin{aligned}\frac{1}{2}\int_0^\infty \frac{u^{\frac{d_\Theta}{2}-1}}{(1+\hat{\eta}^2+u)^s} du =& \frac{1}{2}\int_0^\infty \frac{((1+\hat{\eta}^2)t)^{\frac{d_\Theta}{2}-1}}{(1+\hat{\eta}^2+(1+\hat{\eta}^2)t)^s} 1+\hat{\eta}^2 dt \\ =& \frac{1}{2}\int_0^\infty \frac{((1+\hat{\eta}^2)t)^{\frac{d_\Theta}{2}-1}}{(1+\hat{\eta}^2)(1+t)^s} 1+\hat{\eta}^2 dt \\ =& \frac{1}{2}\int_0^\infty \frac{(1+\hat{\eta}^2)^{\frac{d_\Theta}{2}} t^{\frac{d_\Theta}{2}-1}}{(1+\hat{\eta}^2)^s (1+t)^s} dt \\ =& \frac{1}{2}(1+\hat{\eta}^2)^{\frac{d_\Theta}{2}-s} \int_0^\infty \frac{t^{\frac{d_\Theta}{2}-1}}{(1+t)^s} dt.\end{aligned} \tag{109}$$

The definition of the beta function is $B(a,b) = \int_0^\infty \frac{t^{a-1}}{(1+t)^{a+b}} dt$. Let $a = \frac{d_\Theta}{2}, b = s - \frac{d_\Theta}{2}$, and we have:

$$\int_0^\infty \frac{t^{\frac{d_\Theta}{2}-1}}{(1+t)^s} dt = B(a,b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)} = \frac{\Gamma(\frac{d_\Theta}{2})\Gamma(s-\frac{d_\Theta}{2})}{\Gamma(s)}. \tag{110}$$

Thus, we have:

$$\int_{\mathbb{R}^{d_\Theta \times \mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} \|\eta\|_2^2 d\xi d\eta = \frac{1}{2}\omega_{d_\Theta}\omega_{d_\mathcal{X}} \frac{\Gamma(\frac{d_\Theta}{2})\Gamma(s-\frac{d_\Theta}{2})}{\Gamma(s)} \int_0^\infty (1+\hat{\eta}^2)^{\frac{d_\Theta}{2}-s} \hat{\eta}^{d_\mathcal{X}+1} d\hat{\eta}. \tag{111}$$

Since $\omega_n = \frac{2\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2})}$, we have:

$$\frac{1}{2}\omega_{d_\Theta}\Gamma(\frac{d_\Theta}{2}) = \frac{1}{2}\frac{2\pi^{\frac{d_\Theta}{2}}}{\Gamma(\frac{d_\Theta}{2})}\Gamma(\frac{d_\Theta}{2}) = \pi^{\frac{d_\Theta}{2}}. \tag{112}$$

Thus, we have:

$$\int_{\mathbb{R}^{d_\Theta \times \mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-s} \|\eta\|_2^2 d\xi d\eta = \pi^{\frac{d_\Theta}{2}} \omega_{d_\mathcal{X}} \frac{\Gamma(s-\frac{d_\Theta}{2})}{\Gamma(s)} \int_0^\infty (1+\hat{\eta}^2)^{\frac{d_\Theta}{2}-s} \hat{\eta}^{d_\mathcal{X}+1} d\hat{\eta}. \tag{113}$$

Now, let $u = \hat{\eta}^2$, then $\hat{\eta} = u^{\frac{1}{2}}, d\hat{\eta} = \frac{1}{2}u^{-\frac{1}{2}}du$, and we have:

$$\int_0^\infty (1+\hat{\eta}^2)^{\frac{d_\Theta}{2}-s}\hat{\eta}^{d_\mathcal{X}+1}d\hat{\eta} = \frac{1}{2}\int_0^\infty \frac{u^{\frac{d_\mathcal{X}+1}{2}}u^{-\frac{1}{2}}}{(1+u)^{s-\frac{d_\Theta}{2}}}du = \frac{1}{2}\int_0^\infty \frac{u^{\frac{d_\mathcal{X}}{2}}}{(1+u)^{s-\frac{d_\Theta}{2}}}du. \tag{114}$$

Let $a = \frac{d_\mathcal{X}}{2}+1, b = s-\frac{d_\Theta+d_\mathcal{X}}{2}-1$, then according to the definition of the beta function, we have:

$$\int_0^\infty \frac{u^{\frac{d_\mathcal{X}}{2}}}{(1+u)^{s-\frac{d_\Theta}{2}}}du = \int_0^\infty \frac{u^{a-1}}{(1+u)^{a+b}}du = B(a,b) = \frac{\Gamma(\frac{d_\mathcal{X}}{2}+1)\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{\Gamma(s-\frac{d_\Theta}{2})}. \tag{115}$$

Then, we have:

$$\begin{aligned}\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2 d\xi d\eta =& \pi^{\frac{d_\Theta}{2}}\omega_{d_\mathcal{X}}\frac{\Gamma(s-\frac{d_\Theta}{2})}{\Gamma(s)}\int_0^\infty (1+\hat{\eta}^2)^{\frac{d_\Theta}{2}-s}\hat{\eta}^{d_\mathcal{X}+1}d\hat{\eta}\\ =& \pi^{\frac{d_\Theta}{2}}\omega_{d_\mathcal{X}}\frac{\Gamma(s-\frac{d_\Theta}{2})}{\Gamma(s)}\frac{1}{2}\frac{\Gamma(\frac{d_\mathcal{X}}{2}+1)\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{\Gamma(s-\frac{d_\Theta}{2})}\\ =& \pi^{\frac{d_\Theta}{2}}\omega_{d_\mathcal{X}}\frac{\Gamma(\frac{d_\mathcal{X}}{2}+1)\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{2\Gamma(s)}.\end{aligned} \tag{116}$$

Since $\omega_n = \frac{2\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2})}$, we have $\omega_{d_\mathcal{X}} = \frac{2\pi^{\frac{d_\mathcal{X}}{2}}}{\Gamma(\frac{d_\mathcal{X}}{2})}$. Since the Gamma function satisfies $z\Gamma(z) = \Gamma(z+1)$, we have $\Gamma(\frac{d_\mathcal{X}}{2}+1) = \frac{d_\mathcal{X}}{2}\Gamma(\frac{d_\mathcal{X}}{2})$. Thus, we have:

$$\begin{aligned}\pi^{\frac{d_\Theta}{2}}\omega_{d_\mathcal{X}}\frac{\Gamma(\frac{d_\mathcal{X}}{2}+1)\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{2\Gamma(s)} =& \pi^{\frac{d_\Theta}{2}}\frac{2\pi^{\frac{d_\mathcal{X}}{2}}}{\Gamma(\frac{d_\mathcal{X}}{2})}\frac{\frac{d_\mathcal{X}}{2}\Gamma(\frac{d_\mathcal{X}}{2})\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{2\Gamma(s)}\\ =& \frac{\pi^{\frac{d_{\Theta\times\mathcal{X}}}{2}}d_\mathcal{X}\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{2\Gamma(s)}.\end{aligned} \tag{117}$$

The $\|\eta\|_2^2$ term requires $s > \frac{d_{\Theta\times\mathcal{X}}}{2}+1$ for the Gamma function $\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)$. This condition is satisfied by the chosen $s=3$ with $d_{\Theta\times\mathcal{X}}=2$. As a result, we have that:

$$\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2 d\xi d\eta = \frac{\pi^{\frac{d_{\Theta\times\mathcal{X}}}{2}}d_\mathcal{X}\Gamma(s-\frac{d_\Theta+d_\mathcal{X}}{2}-1)}{2\Gamma(s)}. \tag{118}$$

Since $d_\Theta = 1, d_\mathcal{X} = 1, d_{\Theta\times\mathcal{X}} = 2$ in this experiment, to ensure $\int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2 d\xi d\eta \le 1$, we need to ensure:

$$\pi\frac{\Gamma(s-2)}{2\Gamma(s)} \le 1. \tag{119}$$

When $s=3$, we have $\pi\frac{\Gamma(s-2)}{2\Gamma(s)} = \pi\frac{\Gamma(1)}{2\Gamma(3)}$. Since $\Gamma(1)=1, \Gamma(3)=2$, we have:

$$\pi\frac{\Gamma(s-2)}{2\Gamma(s)} = \pi\frac{\Gamma(1)}{2\Gamma(3)} = \frac{\pi}{4} < 1. \tag{120}$$

Thus, in this experiment, when we set $s = 3$, we have the coefficient $a = \frac{\pi}{4}$:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) + \gamma C(x, \theta, m_1) - \mathcal{L}_\theta(\bar{m}, m_2) - \gamma C(x, \theta, m_2)\right)(m_1 - m_2)(d\theta, dx) \\
=&\eta\|\bar{x}_1 - \bar{x}_2\|_2^2 + \gamma a \int_\Theta \|x_1^*(\theta) - x_2^*(\theta)\|_2^2 \mu(d\theta) \\
\geq&\gamma a \int_\Theta \|x_1^*(\theta) - x_2^*(\theta)\|_2^2 \mu(d\theta) \\
\geq&\frac{\gamma a \|m_1 - m_2\|_{H^{-3}(\Theta\times\mathcal{X})}^2}{\int_{\mathbb{R}^{d_{\Theta\times\mathcal{X}}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-3} \|\eta\|_2^2 d\xi d\eta} \\
=&\gamma\|m_1 - m_2\|_{H^{-3}(\Theta\times\mathcal{X})}^2.
\end{aligned} \tag{121}$$

As a result, after adding the congestion term $\gamma a \int_\Theta \|x_1^*(\theta) - x_2^*(\theta)\|_2^2 \mu(d\theta)$ and setting $s = 3$, the recommendation system game becomes γ-strongly monotone.

H.3 Calculation of Mean-Field Equilibrium

The first optimality condition:

$$\nabla_x \ell = (p - \theta)z - \eta - \eta\bar{x} - \beta(1 + x) = 0. \tag{122}$$

We have:

$$x^*(\theta) = \frac{(p - \theta)z - \eta - \eta\bar{x}}{\beta} - 1. \tag{123}$$

For the mean-field $\bar{x}$, since $\theta \sim \mathcal{U}[0.3p, 0.7p]$:

$$\begin{aligned}
\bar{x} =&\mathbb{E}\left[x^*(\theta)\right] = \mathbb{E}\left[\frac{(p - \theta)z - \eta - \eta\bar{x}}{\beta} - 1\right] \\
=&\frac{(p - \mathbb{E}\left[\theta\right])z - \eta - \eta\bar{x}}{\beta} - 1 \\
=&\frac{(p - 0.5p)z - \eta - \eta\bar{x}}{\beta} - 1.
\end{aligned} \tag{124}$$

Then, we have:

$$\bar{x} = \frac{(p - 0.5p)z}{\eta + \beta} - 1. \tag{125}$$

Then, the equilibrium decision of the player with type θ is:

$$x^*(\theta) = \frac{(p - \theta)z - \eta - \eta\frac{(p-0.5p)z}{\eta+\beta} + \eta}{\beta} - 1 = \frac{(p - \theta)z}{\beta} - \frac{\eta(p - 0.5p)z}{\beta(\eta + \beta)} - 1. \tag{126}$$

If the solution falls outside $[-1, 1]$, the decision is clipped to the boundary, and the mean-field equilibrium is recomputed, analogous to the procedure in Appendix G.

H.4 Calculation of ε-Sensitivity

From the loss function, we have:

$$\left\|\mathcal{L}_\theta(m^{(t-1)}, m) - \mathcal{L}_\theta(m^{(t)}, m)\right\|_{H^3(\Theta\times\mathcal{X})} = \left\|(1 + x)(p - \theta) \cdot (z^{(t-1)} - z^{(t)})\right\|_{H^3(\Theta\times\mathcal{X})}. \tag{127}$$

We have:

$$\varepsilon = \max\left(\varepsilon, \frac{\left\|(1 + x)(p - \theta) \cdot (z^{(t-1)} - z^{(t)})\right\|_{H^3(\Theta\times\mathcal{X})}}{\left\|m^{(t)} - m^{(t-1)}\right\|_{H^{-3}(\Theta\times\mathcal{X})}}\right). \tag{128}$$

We set $f = (1+x)(p-\theta)\cdot(z^{(t-1)} - z^{(t)})$. According to the definition of the Sobolev norm, we have:

$$\|f\|_{H^3(\Theta\times\mathcal{X})} = \left(\sum_{|\beta|\leq 3} \|D^\beta f\|^2_{L^2(\Theta\times\mathcal{X})}\right)^{\frac{1}{2}}, \tag{129}$$

where $D^\beta f$ is a weak β-th partial derivative of f.

For a multi-index $\beta = (\beta_1, \beta_2)$, we have

$$D^\beta f = \frac{\partial^{|\beta|} f}{\partial\theta^{\beta_1}\partial x^{\beta_2}}, \tag{130}$$

where $|\beta| = \beta_1 + \beta_2$.

Then, we have:

$$\begin{aligned}\|f\|_{H^3(\Theta\times\mathcal{X})} &= \left(\sum_{|\beta|\leq 3} \|D^\beta f\|^2_{L^2(\Theta\times\mathcal{X})}\right)^{\frac{1}{2}} \\ &= \left(\|f\|^2_{L^2(\Theta\times\mathcal{X})} + \|\partial_\theta f\|^2_{L^2(\Theta\times\mathcal{X})} + \|\partial_x f\|^2_{L^2(\Theta\times\mathcal{X})} + \|\partial_\theta\partial_x f\|^2_{L^2(\Theta\times\mathcal{X})}\right)^{\frac{1}{2}}.\end{aligned} \tag{131}$$

Since $\mathcal{X} = [-1, 1]$ and $\theta \sim \mathcal{U}[0.3p, 0.7p]$, let $a = 0.3p$ and $b = 0.7p$. Then, for each term, we have:

$$\begin{aligned}\|f\|^2_{L^2(\Theta\times\mathcal{X})} &= (z^{(t-1)} - z^{(t)})^2 \int_{-1}^{1}(1+x)^2 dx \int_a^b (p-\theta)^2 d\theta \\ &= (z^{(t-1)} - z^{(t)})^2 \cdot \frac{8}{3}\cdot\frac{(p-a)^3 - (p-b)^3}{3}, \\ \|\partial_\theta f\|^2_{L^2(\Theta\times\mathcal{X})} &= (z^{(t-1)} - z^{(t)})^2 \int_{-1}^{1}(1+x)^2 dx \int_a^b 1 d\theta \\ &= (z^{(t-1)} - z^{(t)})^2 \cdot \frac{8}{3}\cdot(b-a), \\ \|\partial_x f\|^2_{L^2(\Theta\times\mathcal{X})} &= (z^{(t-1)} - z^{(t)})^2 \int_{-1}^{1}1 dx \int_a^b (p-\theta)^2 d\theta \\ &= (z^{(t-1)} - z^{(t)})^2 \cdot 2\cdot\frac{(p-a)^3 - (p-b)^3}{3}, \\ \|\partial_\theta\partial_x f\|^2_{L^2(\Theta\times\mathcal{X})} &= (z^{(t-1)} - z^{(t)})^2 \int_{-1}^{1}1 dx \int_a^b 1 d\theta \\ &= (z^{(t-1)} - z^{(t)})^2 \cdot 2\cdot(b-a).\end{aligned} \tag{132}$$

From the monotonicity calculation, we have:

$$\begin{aligned}&\|m^{(t)} - m^{(t-1)}\|^2_{H^{-3}(\Theta\times\mathcal{X})} \\ =& \int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} \left(1 + \|\xi\|_2^2 + \|\eta\|_2^2\right)^{-3} \left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x^{*(t)}(\theta)} - e^{-i\eta x^{*(t-1)}(\theta)}\right)\mu(d\theta)\right|^2 d\xi d\eta.\end{aligned} \tag{133}$$

I Additional Details of Ride Share Market

I.1 Implementation Details

In a ride share market, players choose order volume adjustments on Uber and Lyft to maximize utility. We use ride share order data (Narang et al., 2023)[5] (14 days of order records across 11 locations

[5] Available at https://www.kaggle.com/datasets/brllrb/uber-and-lyft-dataset-boston-ma.

in Boston), from which we extract 5 price levels $\{10, 15, 20, 25, 30\}$, each treated as an independent mean field game instance with demand aggregated by averaging across all locations and days. Each player has a type $\theta \sim \mathrm{Beta}(10, 10)$ representing their preference for Uber, with $1 - \theta$ being the preference for Lyft. Specifically, each player minimizes the following loss:

$$\min_{x \in \mathcal{X}} \frac{\zeta}{2}\|x\|^2 + \hat{p}^\top(1+x) - \eta t^\top(1+x) + \beta \mathbf{1}^\top(1+x)\mathbf{1}^\top(1+\bar{x}), \tag{134}$$

where $x \in \mathcal{X} = [-1, 1]^2$ is the percentage adjustment in order volume on the two platforms, $\zeta = 0.01$ is the adjustment cost coefficient, $\hat{p} = [\hat{p}_1, \hat{p}_2]^\top$ is the normalized price with $\hat{p}_i = \frac{p_i}{p_{0,i} \cdot 0.1}$ where p_i is the raw price of platform i drawn from $\mathcal{D}_i(m)$ and $p_{0,i}$ is the baseline price of platform i, $t = [\theta, 1-\theta]^\top$ is the preference vector, $\eta = 30$ is the preference weight, $\bar{x} = \int_{\Theta \times \mathcal{X}} x m(d\theta, dx)$ is the average adjustment vector, and $\beta \mathbf{1}^\top(1+x)\mathbf{1}^\top(1+\bar{x})$ is the traffic congestion cost. The congestion coefficient β for each price level is set proportionally to its cumulative demand share: $\beta_i = \beta_0 \cdot \frac{\sum_{j \le i} d_j}{\sum_j d_j}$, where $\beta_0 = 0.1$ is the base congestion coefficient and d_j is the total demand (Uber + Lyft) at price level j. This design reflects that higher price levels are associated with longer travel distances and thus heavier traffic congestion, so the congestion cost should increase with the price level. The data distribution map reflects that collective order adjustments affect prices:

$$\begin{aligned} p_1 &\sim \mathcal{D}_1(m) = p_{0,1}\big(1 + (1 + 0.4\kappa)\bar{x}_1|\bar{x}_1|\big) + w_1, \\ p_2 &\sim \mathcal{D}_2(m) = p_{0,2}\big(1 + (1 + 0.6\kappa)\bar{x}_2^3\big) + w_2, \end{aligned} \tag{135}$$

where $p_{0,1}$ and $p_{0,2}$ are the baseline prices, κ is the performative effect strength, and $w_1, w_2 \sim \mathcal{N}(0, 1)$ are noise terms. Uber's price responds quadratically with respect to the collective order volume adjustment, while Lyft's responds cubically. We set $\kappa \in \{0.25, 0.5, 0.75, 1.0\}$.

We draw $N_\theta = 500$ type samples from $\mathrm{Beta}(10, 10)$ and $N_\omega = 1000$ frequency samples. The H^{-s} norm is computed via 3D Monte Carlo Fourier integration ($d_\Theta = 1$, $d_\mathcal{X} = 2$, $d_{\Theta \times \mathcal{X}} = 3$) with multivariate Cauchy importance sampling. The convergence threshold is set to $\rho = 10^{-1}$. All experiments are repeated with 10 different random seeds (42–51), and the reported results are averaged over these 10 runs. All experiments were performed on a server running Oracle Linux Server 8.10 with an AMD EPYC 7742 64-Core CPU (256 threads) and 1 TB of RAM, without GPU acceleration. The code is implemented in Python 3.12.10, with key libraries including NumPy 2.5.0, SciPy 1.18.0, and Matplotlib 3.11.0.

I.2 Monotonicity

In the following derivations, p denotes the normalized price $\hat{p}$ defined in the Implementation Details.

According to the definition, we have:

$$\begin{aligned} &\int_{\Theta \times \mathcal{X}} \left(\mathcal{L}_\theta(\bar{m}, m_1) - \mathcal{L}_\theta(\bar{m}, m_2)\right)(m_1 - m_2)(d\theta, dx) \\ =&\int_{\Theta \times \mathcal{X}} \left(\left(\frac{\zeta}{2}\|x\|_2^2 + (p - \eta t)^T(1+x) + \beta \mathbf{1}^T(1+x)\mathbf{1}^T(1+\bar{x}_1)\right) - \left(\frac{\zeta}{2}\|x\|_2^2 + (p - \eta t)^T(1+x) + \beta \mathbf{1}^T(1+x)\mathbf{1}^T(1+\bar{x}_2)\right)\right)(m_1 - m_2)\,(d\theta, dx) \\ =&\int_{\Theta \times \mathcal{X}} \beta \mathbf{1}^T(1+x)\mathbf{1}^T(\bar{x}_1 - \bar{x}_2)\,(m_1 - m_2)\,(d\theta, dx) \\ =&\beta \mathbf{1}^T(\bar{x}_1 - \bar{x}_2) \int_{\Theta \times \mathcal{X}} \mathbf{1}^T(1+x)\,(m_1 - m_2)\,(d\theta, dx) \\ =&\beta \mathbf{1}^T(\bar{x}_1 - \bar{x}_2)\mathbf{1}^T(\bar{x}_1 - \bar{x}_2) \\ =&\beta \|\mathbf{1}^T(\bar{x}_1 - \bar{x}_2)\|_2^2. \end{aligned} \tag{136}$$

Since the decision distribution given type θ in the mean-field equilibrium is a Dirac measure, every measure $m \in \mathcal{M}_{RPMFG}$ satisfies that, for all measurable sets $A \subset \Theta$ and $B \subset \mathcal{X}$, $m(A \times B) = \int_A \delta_{x^*(\theta)}(B)\mu(d\theta)$, where μ is the fixed probability measure over the type space Θ, and $\delta_{x^*(\theta)}$ is the Dirac measure on $\mathcal{X}$ at the point $x^*(\theta)$ for type θ.

We consider the congestion term $C(x, \theta, m) = a \cdot x \cdot \int_\mathcal{X} x m(dx|\theta)$, where a is a coefficient. Since the decision distribution in mean-field equilibrium is a Dirac measure, $\int_\mathcal{X} x m(dx|\theta) = x^*(\theta)$ and

$C(x,\theta,m) = a\cdot x\cdot\int_{\mathcal{X}} xm(dx|\theta) = a\cdot x\cdot x^*(\theta)$. In this case, the monotonicity of the game becomes:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m},m_1)+\gamma C(x,\theta,m_1)-\mathcal{L}_\theta(\bar{m},m_2)-\gamma C(x,\theta,m_2)\right)(m_1-m_2)(d\theta,dx)\\
=&\beta\|\mathbf{1}^T(\bar{x}_1-\bar{x}_2)\|_2^2+\gamma\cdot a\cdot\int_\Theta \|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)\\
=&\beta\|\mathbf{1}^T(\bar{x}_1-\bar{x}_2)\|_2^2+\gamma\cdot a\cdot\int_\Theta \left\|\int_{\mathcal{X}} x'\left(\delta_{x_1^*(\theta)}(dx')-\delta_{x_2^*(\theta)}\right)(dx')\right\|_2^2\mu(d\theta).
\end{aligned} \tag{137}$$

Since every measure $m\in\mathcal{M}_{RPMFG}$ satisfies that, for all measurable sets $A\subset\Theta$ and $B\subset\mathcal{X}$, $m(A\times B)=\int_A \delta_{x^*(\theta)}(B)\mu(d\theta)$, $\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}$ can be calculated using Fourier transformation (Cerdà, 2010). The detailed calculation is similar to that in the Quantitative Trading Game, as shown in Appendix G. The Gamma function in the closed-form expression requires $s>\frac{d_{\Theta\times\mathcal{X}}}{2}+1=\frac{5}{2}$, which is satisfied by the chosen $s=4$.

Since $d_\Theta=1, d_{\mathcal{X}}=2, d_{\Theta\times\mathcal{X}}=3$ in this experiment, to ensure $\int_{\mathbb{R}^{d_{\Theta\times\mathcal{X}}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-s}\|\eta\|_2^2 d\xi d\eta\le 1$, we need to ensure:

$$\frac{\pi^{\frac{3}{2}}2\Gamma(s-\frac{1+2}{2}-1)}{2\Gamma(s)}=\pi^{\frac{3}{2}}\frac{\Gamma(s-\frac{5}{2})}{\Gamma(s)}\le 1. \tag{138}$$

When $s=4$, we have $\pi^{\frac{3}{2}}\frac{\Gamma(s-\frac{5}{2})}{\Gamma(s)}=\pi^{\frac{3}{2}}\frac{\Gamma(\frac{3}{2})}{\Gamma(4)}$. Since $\Gamma(\frac{3}{2})=\frac{1}{2}\pi^{\frac{1}{2}}, \Gamma(4)=6$, we have:

$$\pi^{\frac{3}{2}}\frac{\Gamma(s-\frac{5}{2})}{\Gamma(s)}=\pi^{\frac{3}{2}}\frac{\Gamma(\frac{3}{2})}{\Gamma(4)}=\frac{\pi^2}{12}<1. \tag{139}$$

Thus, in this experiment, when we set $s=4$, we have $a=\frac{\pi^2}{12}$:

$$\begin{aligned}
&\int_{\Theta\times\mathcal{X}} \left(\mathcal{L}_\theta(\bar{m},m_1)+\gamma C(x,\theta,m_1)-\mathcal{L}_\theta(\bar{m},m_2)-\gamma C(x,\theta,m_2)\right)(m_1-m_2)(d\theta,dx)\\
=&\beta\|\mathbf{1}^T(\bar{x}_1-\bar{x}_2)\|_2^2+\gamma\cdot a\cdot\int_\Theta \|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)\\
\ge&\gamma\cdot a\cdot\int_\Theta \|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)\\
\ge&\frac{\gamma\cdot a\cdot\|m_1-m_2\|^2_{H^{-s}(\Theta\times\mathcal{X})}}{\int_{\mathbb{R}^{d_{\Theta\times\mathcal{X}}}}\left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-4}\|\eta\|_2^2 d\xi d\eta}\\
=&\gamma\|m_1-m_2\|^2_{H^{-4}(\Theta\times\mathcal{X})}.
\end{aligned} \tag{140}$$

As a result, after adding the congestion term $\gamma\cdot a\cdot\int_\Theta\|x_1^*(\theta)-x_2^*(\theta)\|_2^2\mu(d\theta)$ and setting $s=4$, the ride-share game becomes γ-strongly monotone.

I.3 Calculation of Mean-Field Equilibrium

First-order optimality:

$$\nabla_x\ell=\zeta x+p-\eta t+\beta\mathbf{1}^T(\mathbf{1}+\bar{x})\mathbf{1}. \tag{141}$$

Then:

$$x^*(\theta,\bar{x})=\frac{\eta t-p-\beta\mathbf{1}^T(\mathbf{1}+\bar{x})\mathbf{1}}{\zeta}. \tag{142}$$

Mean-field:

$$\mathbb{E}[x^*(\theta,\bar{x})]=\mathbb{E}\left[\frac{\eta t-p-\beta\mathbf{1}^T(\mathbf{1}+\bar{x})\mathbf{1}}{\zeta}\right]=\frac{\frac{\eta}{2}\cdot\mathbf{1}-p-\beta\mathbf{1}^T(\mathbf{1}+\mathbb{E}[x^*(\theta,\bar{x})])\mathbf{1}}{\zeta}. \tag{143}$$

Then, we have:

$$\begin{aligned}
&\zeta\mathbb{E}[x^*(\theta,\bar{x})] = \frac{\eta}{2}\cdot\mathbf{1} - p - \beta\mathbf{1}^T\mathbf{1}\cdot\mathbf{1} - \beta\mathbf{1}^T\mathbb{E}[x^*(\theta,\bar{x})]\mathbf{1} \\
\Longrightarrow &(\zeta I + \beta\mathbf{1}\mathbf{1}^T)\mathbb{E}[x^*(\theta,\bar{x})] = \frac{\eta}{2}\mathbf{1} - p - \beta\mathbf{1}^T\mathbf{1}\cdot\mathbf{1} \\
\Longrightarrow &\mathbb{E}[x^*(\theta,\bar{x})] = \left(\zeta I + \beta\mathbf{1}\mathbf{1}^T\right)^{-1}\left(\frac{\eta}{2}\mathbf{1} - p - \beta\mathbf{1}^T\mathbf{1}\cdot\mathbf{1}\right).
\end{aligned} \tag{144}$$

Then, for each player, the equilibrium decision is:

$$x^*(\theta) = \frac{\eta t - p - \beta\mathbf{1}^T\left(\mathbf{1} + \left(\zeta I + \beta\mathbf{1}\mathbf{1}^T\right)^{-1}\left(\frac{\eta}{2}\mathbf{1} - p - \beta\mathbf{1}^T\mathbf{1}\cdot\mathbf{1}\right)\right)\mathbf{1}}{\zeta}. \tag{145}$$

If the solution falls outside $[-1,1]^2$, each coordinate is clipped to the boundary, and the mean-field equilibrium is recomputed, analogous to the procedure in Appendix G.

I.4 Calculation of ε-Sensitivity

From the loss function, we have:

$$\|\mathcal{L}_\theta(m_1,m) - \mathcal{L}_\theta(m_2,m)\|_{H^s(\Theta\times\mathcal{X})} = \|(p(m_1)-p(m_2))(1+x)\|_{H^s(\Theta\times\mathcal{X})}. \tag{146}$$

Since we set $s=4$, we have:

$$\varepsilon = \max\left(\varepsilon, \frac{\left\|(p^{(t)}-p^{(t-1)})(1+x)\right\|_{H^4(\Theta\times\mathcal{X})}}{\left\|m^{(t)}-m^{(t-1)}\right\|_{H^{-4}(\Theta\times\mathcal{X})}}\right). \tag{147}$$

We set $f = (p^{(t)}-p^{(t-1)})(1+x)$. According to the definition of the Sobolev norm, we have

$$\|f\|_{H^4(\Theta\times\mathcal{X})} = \left(\sum_{|\beta|\le 4}\|D^\beta f\|^2_{L^2(\Theta\times\mathcal{X})}\right)^{\frac{1}{2}}, \tag{148}$$

where $D^\beta f$ is a weak β-th partial derivative of f.

For a multi-index $\beta = \beta_0,\beta_1,\beta_2$, we have

$$D^\beta f = \frac{\partial^{|\beta|} f}{\partial\theta^{\beta_0}\partial x_1^{\beta_1}\partial x_2^{\beta_2}}. \tag{149}$$

Since f is linear in x and independent of θ, all derivatives of order $|\beta|\ge 2$ are 0. The only non-zero terms are those with $|\beta|=0$ and $|\beta|=1$. Then, we have:

$$\begin{aligned}
\|f\|_{H^4(\Theta\times\mathcal{X})} &= \left(\sum_{|\beta|\le 4}\|D^\beta f\|^2_{L^2(\Theta\times\mathcal{X})}\right)^{\frac{1}{2}} \\
&= \left(\|f\|^2_{L^2(\Theta\times\mathcal{X})} + \|\partial_{x_1} f\|^2_{L^2(\Theta\times\mathcal{X})} + \|\partial_{x_2} f\|^2_{L^2(\Theta\times\mathcal{X})}\right)^{\frac{1}{2}}.
\end{aligned} \tag{150}$$

For each term, we have:

$$
\begin{aligned}
\|f\|^2_{L^2(\Theta\times\mathcal{X})} &= \int_0^1\int_{-1}^1\int_{-1}^1 \left((p^{(t)}-p^{(t-1)})(1+x)\right)^2 d\theta dx_1 dx_2 \\
&= \int_{-1}^1\int_{-1}^1 \left((p^{(t)}-p^{(t-1)})(1+x)\right)^2 dx_1 dx_2 \\
&= \int_{-1}^1\int_{-1}^1 \left((p_1^{(t)}-p_1^{(t-1)})(1+x_1)+(p_2^{(t)}-p_2^{(t-1)})(1+x_2)\right)^2 dx_1 dx_2 \\
&= \int_{-1}^1\int_{-1}^1 \left((p_1^{(t)}-p_1^{(t-1)})^2(1+x_1)^2+(p_2^{(t)}-p_2^{(t-1)})^2(1+x_2)^2+2(p_1^{(t)}-p_1^{(t-1)})(1+x_1)(p_2^{(t)}-p_2^{(t-1)})(1+x_2)\right) dx_1 dx_2 \\
&= \frac{16}{3}(p_1^{(t)}-p_1^{(t-1)})^2+\frac{16}{3}(p_2^{(t)}-p_2^{(t-1)})^2+8(p_1^{(t)}-p_1^{(t-1)})(p_2^{(t)}-p_2^{(t-1)}), \\
\|\partial_{x_1} f\|^2_{L^2(\Theta\times\mathcal{X})} &= \int_0^1\int_{-1}^1\int_{-1}^1 (p_1^{(t)}-p_1^{(t-1)})^2 d\theta dx_1 dx_2 \\
&= 4(p_1^{(t)}-p_1^{(t-1)})^2, \\
\|\partial_{x_2} f\|^2_{L^2(\Theta\times\mathcal{X})} &= \int_0^1\int_{-1}^1\int_{-1}^1 (p_2^{(t)}-p_2^{(t-1)})^2 d\theta dx_1 dx_2 \\
&= 4(p_2^{(t)}-p_2^{(t-1)})^2.
\end{aligned}
\tag{151}
$$

Then, we have:

$$
\begin{aligned}
&\left\|(p^{(t)}-p^{(t-1)})(1+x)\right\|_{H^4(\Theta\times\mathcal{X})} \\
&= \left(\frac{28}{3}\left((p_1^{(t)}-p_1^{(t-1)})^2+(p_2^{(t)}-p_2^{(t-1)})^2\right)+8(p_1^{(t)}-p_1^{(t-1)})(p_2^{(t)}-p_2^{(t-1)})\right)^{\frac{1}{2}}.
\end{aligned}
\tag{152}
$$

From the monotonicity calculation, we have:

$$
\begin{aligned}
&\|m^{(t)}-m^{(t-1)}\|^2_{H^{-4}(\Theta\times\mathcal{X})} \\
&= \int_{\mathbb{R}^{d_\Theta\times\mathcal{X}}} \left(1+\|\xi\|_2^2+\|\eta\|_2^2\right)^{-4} \left|\int_\Theta e^{-i\xi\theta}\left(e^{-i\eta x^{*(t)}(\theta)}-e^{-i\eta x^{*(t-1)}(\theta)}\right)\mu(d\theta)\right|^2 d\xi d\eta.
\end{aligned}
\tag{153}
$$

J Additional Experimental Results

J.1 $\|m^{(t)}-m^{(t-1)}\|_{H^{-s}}$ across Iterations

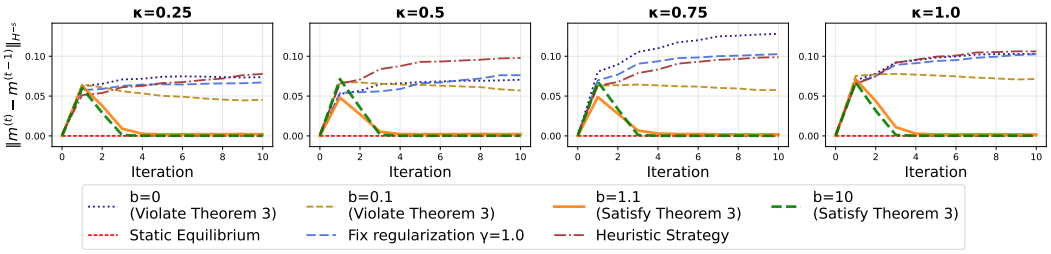

Figure 2: $\|m^{(t)}-m^{(t-1)}\|_{H^{-s}}$ in quantitative trading with our algorithm ($b=0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma=1$, Heuristic Strategy) over 10 trials.

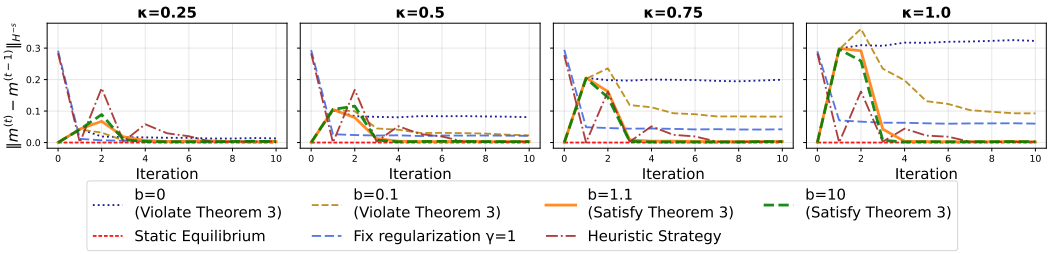

Figure 3: $\|m^{(t)} - m^{(t-1)}\|_{H^{-s}}$ in recommendation systems with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

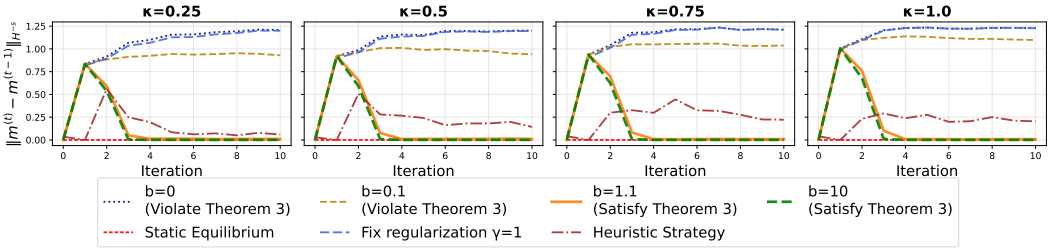

Figure 4: $\|m^{(t)} - m^{(t-1)}\|_{H^{-s}}$ in ride share market with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

J.2 Average Decision across Iterations

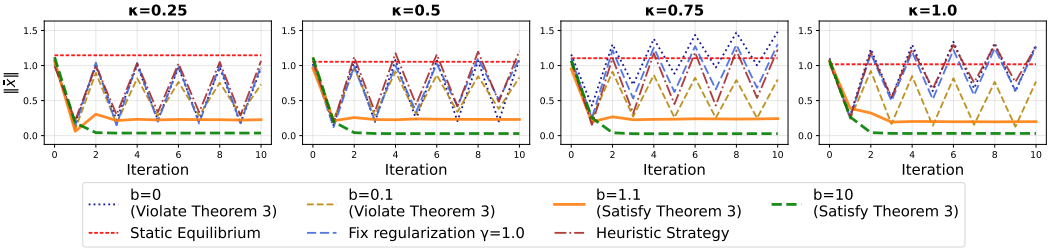

Figure 5: L_2-norm of average decision $\|\bar{x}^{(t)}\|_2$ in quantitative trading with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

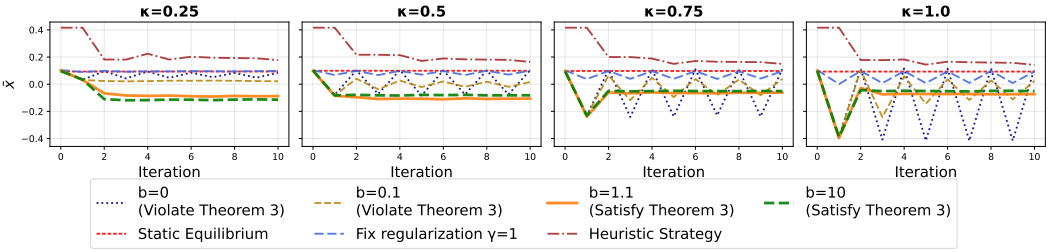

Figure 6: Average price adjustment $\bar{x}^{(t)}$ in recommendation systems with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

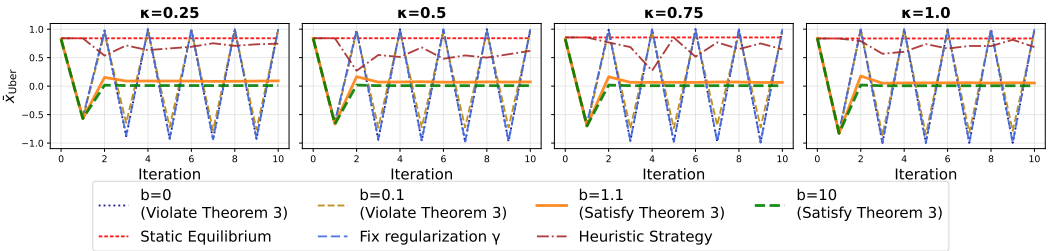

Figure 7: Average order adjustment for Uber $\bar{x}^{(t)}_{\text{Uber}}$ in ride share market with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

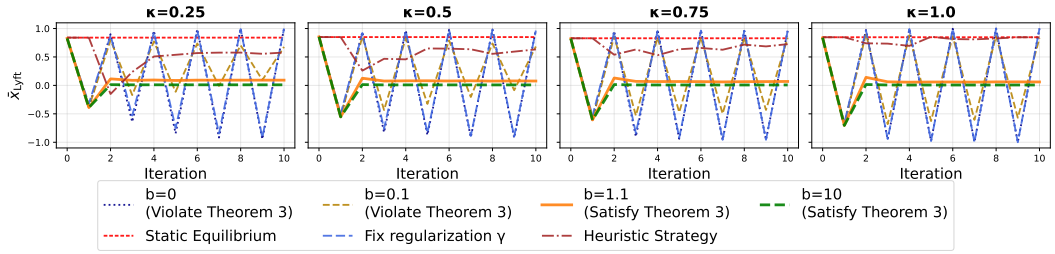

Figure 8: Average order adjustment for Lyft $\bar{x}^{(t)}_{\text{Lyft}}$ in ride share market with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

J.3 Average Payoff with Standard Deviation across Iterations

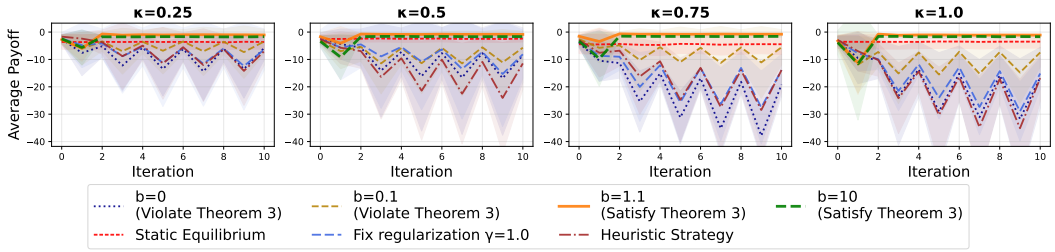

Figure 9: Average payoff with standard deviation in quantitative trading with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

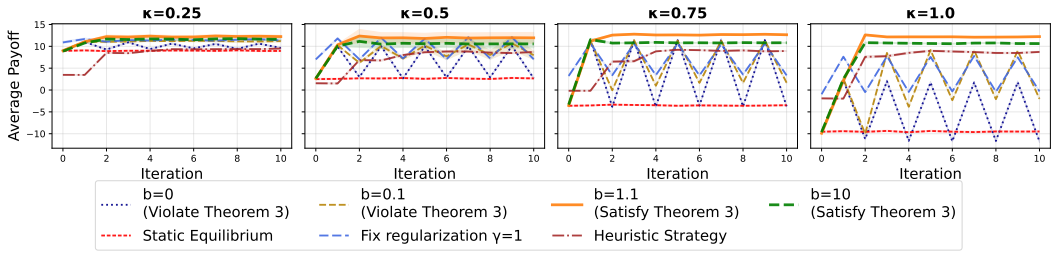

Figure 10: Average payoff with standard deviation in recommendation systems with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

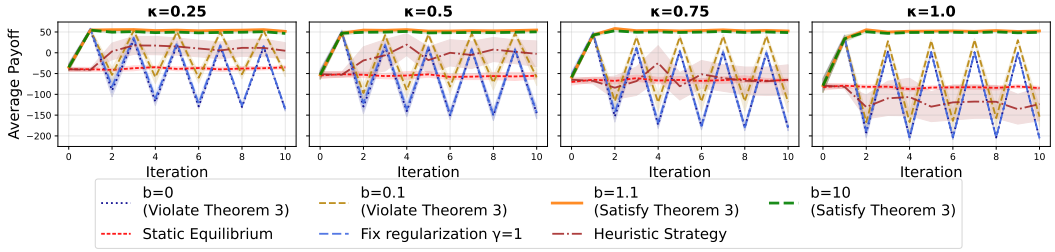

Figure 11: Average payoff with standard deviation in ride share market with our algorithm ($b = 0, 0.1, 1.1, 10$) and three baselines (Static Equilibrium, Fixed Regularization $\gamma = 1$, Heuristic Strategy) over 10 trials.

Table 1: Wilcoxon signed-rank test p-values (two-sided, Bonferroni-corrected). Each comparison uses 12 paired settings (3 games $\times$ 4 values of κ).

	$b = 0$	$b = 0.1$	$b = 1.1$	$b = 10$	Static Equilibrium	Fixed Regularization	Heuristic Strategy
$b = 1.1$	0.0005	0.0005	—	0.0005	0.0005	0.0005	0.0005
$b = 10$	0.0005	0.0005	0.0005	—	0.0005	0.0005	0.0005

J.4 Statistical Significance

We conduct pairwise two-sided Wilcoxon signed-rank tests (Wilcoxon, 1945) using the payoff at the final iteration of each method in 12 settings (3 games $\times$ 4 values of κ) as paired observations. For our algorithm with $b = 1.1$ and $b = 10$, we pair their 12 final payoffs with the corresponding 12 final payoffs of each comparison method, yielding one p-value per comparison. We set the significance level to 0.05 and apply Bonferroni correction for the 12 pairwise comparisons. These 12 pairwise comparisons in the Bonferroni correction consist of 2 choices of b ($b = 1.1$ and $b = 10$) in Algorithm 1 compared against 6 other methods, giving $\alpha_{\text{Bonf}} = \frac{0.05}{12} \approx 0.00417$.

All p-values in Table 1 are $p = 0.0005 < \alpha_{\text{Bonf}} \approx 0.00417$, indicating statistical significance at the 0.05 significance level after Bonferroni correction. The value $p = 0.0005$ is the minimum achievable p-value for a two-sided Wilcoxon test with 12 paired observations ($p = 2 \times 0.5^{12} \approx 0.0005$), corresponding to the case where all paired differences share the same sign. The Wilcoxon test operates on the ranks of absolute differences. Since all 12 pairwise differences share the same sign in every comparison we report, the test outcome is driven by this sign consistency rather than by the magnitude of the ranks. These results provide evidence that Algorithm 1 with $b > 1$ ($b = 1.1, 10$) performs significantly differently from Algorithm 1 with $b < 1$ ($b = 0, 0.1$) and from all three baselines. The difference between Algorithm 1 with $b = 1.1$ and $b = 10$ is also significant.